Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Real-Time Machine Learning Control of a Hall Thruster Discharge Plasma

Ajay Krishnan

ajaykrishnan@gatech.edu

Georgia Institute of Technology

Bogdan Vlahov

Georgia Institute of Technology

Jason Gibson

Georgia Institute of Technology

Dan Lev

Georgia Institute of Technology

Evangelos Theodorou

Georgia Institute of Technology

Mitchell Walker

Georgia Institute of Technology

Maryam Saeedifard

Georgia Institute of Technology

Research Article

Keywords:

Posted Date: April 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9325578/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Real-Time Machine Learning Control of a Hall Thruster Discharge Plasma

Ajay Krishnan*, Bogdan Vlahov†, Jason Gibson‡, Dan Lev§, Evangelos Theodorou¶, Mitchell Walker‖ and Maryam Saeedifard**

Georgia Institute of Technology, Atlanta, Georgia, 30332, USA

Facility effects are known to alter Hall thruster performance between ground and space. Furthermore, discharge current oscillations are linked to impacting both the lifetime and thrust of Hall thrusters. In this work, we present a machine learning–based control framework designed for nonlinear, time-varying systems with hysteresis that can mitigate breathing mode oscillations and adapt to environmental changes. Unlike reactive approaches such as PID control, our method is predictive and flexible across a range of control objectives. Specifically, we employ echo state networks (ESNs), a class of recurrent neural networks well-suited for time-series prediction. To train and test the model, we apply a diverse set of discharge voltage perturbations—including sine, triangle, square, ramp, chirp, bandlimited gaussian noise, pseudorandom binary sequence, and PID—on a 6-kW H6 Hall thruster operated at 400 V as well as 250 V at 4.3 A, and measure the resulting discharge current as a part of a process known as system identification. Using extended convergent cross mapping, we confirm causal relationships between the applied perturbations and discharge current response. The ESN achieves 90% prediction accuracy on previously unseen time-series data based on variance. We then integrate the ESN into a nonlinear model predictive control (NMPC) framework, creating a machine learning controller, to compute optimal control trajectories that suppress discharge oscillations. NMPC simulations indicate that discharge current oscillations can be reduced by over 90%. Finally, we employ a deployment strategy in which the learned control policy is embedded as a multilayer perceptron on a field-programmable gate array, enabling real-time experimental implementation. The machine learning controller reduces oscillations by 1.33% while PID increases oscillations by 4.17% at the 250 V operating point. The machine learning (ML) controller has the potential to be applicable to other operating points by following a standard procedure of system identification and control optimization. Online and adaptive ML control can be used in the future to improve oscillation reduction.

I. Nomenclature

HPEPL	=	High Power Electric Propulsion Laboratory	ODE	=	Ordinary Differential Equation
HLS	=	High Level Synthesis	PDE	=	Partial Differential Equation
PPU	=	Power Processing Unit	HET	=	Hall Effect Thruster
VTF	=	Vacuum Test Facility	MFC	=	Mass Flow Controller
BM	=	Breathing Mode	FFT	=	Fast Fourier Transform
STFT	=	Short-Time Fourier Transform	CWT	=	Continuous Wavelet Transform
EMD	=	Empirical Mode Decomposition	ML	=	Machine Learning

*Corresponding Author, Graduate Research Assistant, School of Electrical and Computer Engineering, ajaykrishnan@gatech.edu

†Graduate Research Assistant, Institute for Robotics and Intelligent Machines, bvlahov3@gatech.edu

‡Graduate Research Assistant, Institute for Robotics and Intelligent Machines, jgibson37@gatech.edu

§Research Engineer, School of Aerospace Engineering, dan.lev@gatech.edu

¶Associate Professor, School of Aerospace Engineering, evangelos.theodorou@gatech.edu

‖Chair, School of Aerospace Engineering, mitchell.walker@coe.gatech.edu

**Professor, School of Electrical and Computer Engineering, maryam@ece.gatech.edu

NN	=	Neural Network	RNN	=	Recurrent Neural Network
ESN	=	Echo State Network	RMSE	=	Root Mean Square Deviation
NRMSE	=	Normalized Root Mean Square Error	MLP	=	Multilayer Perceptron
PID	=	Proportional-Integral-Derivative	DC	=	Direct Current
AC	=	Alternating Current	ILC	=	Iterative Learning Control
NMPC	=	Nonlinear Model Predictive Control	IPOPT	=	Interior Point Optimizer
SQP	=	Sequential Quadratic Programming	AM	=	Amplitude Modulation
FM	=	Frequency Modulation	PM	=	Phase Modulation
eCCM	=	Extended Convergent Cross Mapping	EKF	=	Extended Kalman Filter
ADC	=	Analog-to-Digital Converter	FPGA	=	Field-Programmable Gate Array
DAC	=	Digital-to-Analog Converter	GPU	=	Graphics Processing Unit
CPU	=	Central Processing Unit	LSTM	=	Long Short-Term Memory
RLS	=	Recursive Least Squares	V_P	=	PPU Voltage
RMS	=	Root Mean Square	PP	=	Peak-to-Peak
0D	=	Zero-Dimensional	EMI	=	Electromagnetic Interference
IVB	=	Current, Voltage, Magnetic Field	RLC	=	Resistor, Inductor, Capacitor
LTI	=	Linear Time-Invariant	LTV	=	Linear Time-Varying
NLTI	=	Nonlinear Time-Invariant	NLTV	=	Nonlinear Time-Varying
PRBS	=	Pseudorandom Binary Sequence	SNR	=	Signal-to-Noise Ratio
RF	=	Filter Resistance	CF	=	Filter Capacitance
RH	=	Harness Resistance	LH	=	Harness Capacitance
LH	=	Harness Inductance	I_D	=	Discharge Current
VM	=	Modulation Voltage	e	=	Electron Charge
A_C	=	Channel Cross-Sectional Area	M_i	=	Ion Mass
R_Δ	=	Channel Width	ζ_{ion}	=	Ionization Rate Coefficient
L_{ch}	=	Channel Length	N_{int}	=	Injected Neutral Number Density at Anode
ε_w	=	Electron Energy Loss to Wall	ν_w	=	Electron Wall Collision Frequency
χ	=	Ionization Cost	ε_{ion}	=	Ionization Energy Loss
$U_{i,w}$	=	Ion Acoustic Speed	E	=	Electric Field
V_{CF}	=	Filter Capacitor Voltage	I_{LH}	=	Harness Inductor Current
N_i	=	Ion Number Density	N_n	=	Neutral Number Density
U_i	=	Ion Velocity	U_e	=	Electron Velocity
T_e	=	Electron Temperature	$\hat{h}(n)$	=	Intermediate Reservoir Neuron Activations
$h(n)$	=	Reservoir Neuron Activations	$x(n)$	=	Training Input Signal
$y(n)$	=	Output	W_{in}	=	Input Weight Matrix
W_{out}	=	Output Weight Matrix	W_{rec}	=	Recurrent Weight Matrix
W_{fb}	=	Feedback Connections	b_{rec}	=	Recurrent Bias
b_{out}	=	Output Bias	α	=	Leak Rate
Y_{target}	=	Target Output	β	=	Regularization Term
I	=	Identity Matrix	Y_{pred}	=	Predicted Output
K_P	=	Proportional Gain	K_I	=	Integral Gain
K_D	=	Derivative Gain	V_{AC}	=	AC Voltage Component
I_{AC}	=	AC Current Component	SSF	=	Single-Step Forced

SPSA = Simultaneous Perturbation Stochastic Approximation

HITL = Hardware In The Loop

II. Introduction

A HET is an electrostatic plasma device that produces thrust via the acceleration of ions [1, 2]. Developed in the 1960s, HETs have become increasingly adopted for satellite propulsion due to their mass efficiency compared to chemical propulsion systems [2]. To expand their use in satellites and deep space missions, ongoing advancements aim to improve thrust, specific impulse, efficiency, and thruster lifetime. These developments rely on ground-based testing in dedicated vacuum facilities designed to replicate the space environment as closely as possible. However, space conditions cannot be fully reproduced on the ground, resulting in deviations in measured performance and plasma behavior [3].

Vacuum chambers introduce what are referred to as "facility effects" that change thruster performance on the ground compared to that in space [3–6]. Facility effects occur because ground facility pumping systems cannot reach the low-pressure levels as in space [7]. Additionally, the vacuum chambers are made out of metal, creating current paths that otherwise do not exist in the space environment [3–6]. Lastly, sputtering, originating from the chamber walls back onto and into the thruster, may interfere with the discharge and affect HET lifetime [3]. Understanding and mitigating such facility effects is therefore essential for accurate HET performance characterization.

Discharge current oscillations are another critical factor influencing HET operation [8, 9]. These oscillations originate from plasma instabilities, with the dominant mode being the BM [10]. The BM arises from predator–prey interactions between neutrals and electrons, producing variations in plasma density within the discharge chamber [10]. It can be described using conservation equations for mass, momentum, and energy, which enable modeling of the discharge plasma as an electrical load [11]. Typically occurring in the tens of kHz range, BM oscillations can drive current amplitudes exceeding 100% of the mean [10]. Such oscillatory dynamics degrade thruster performance, contribute to wall erosion, and generate EMI, motivating recent efforts to develop active control strategies [12–15].

HETs are nonlinear, spatially-varying, time-varying, and causal systems that exhibit hysteresis [16, 17]. Plasma diagnostics show spatial dependence, while spatially-averaged data from electrical diagnostics reveal time-varying and hysteretic nature. Nonlinearity of the discharge plasma has been demonstrated by using small-signal linearization to show that discharge impedance is constant for small perturbations but begins to change for larger perturbations [18]. The discharge plasma is, however, LTV and not LTI for smaller perturbations, which can be observed when not using a time constant much longer than the breathing mode period that averages out the time-varying behavior. Analytical techniques for analyzing nonlinear, spatially-varying, and time-varying behavior require the use of algorithms beyond the FFT, which applies only for signals produced by LTI systems. The STFT and CWT may be used for LTV systems, while the EMD and similar methods, such as the IMFogram, may be used for NLTI and NLTV systems [16]. It is important to understand the dynamics of the HET so that applicable control strategies can be applied to handle the dynamics.

While IVB maps can identify operating points where oscillations are lower in amplitude, it is sometimes still necessary to seek to reduce the amplitude even further, and it may even be the case where there is a desire to operate a thruster at a more oscillatory point for performance reasons [19]. A simple case of HET control involves tuning an RLC network placed in the discharge circuit that can be used to reduce oscillations [20, 21]. However, the RLC network cannot be changed at a fast rate and is static. Mass flow rate control can also only be accomplished at slow time scales [22]. Magnetic field control is difficult to achieve in real-time due to the resulting back EMF from changing magnetic fields at a fast rate, and thus it is better to tune to optimal values without or with slow variations [20]. Past research has explored perturbing discharge voltage with PID control and phase synchronization of discharge voltage and discharge current [12, 13, 23]. Due to limitations in using PID for nonlinear time-varying loads such as a HET, some researchers have investigated using ML to control a HET, such as ILC [22, 24]. Existing discharge plasma models are not fully predictive, whereas ML has the potential to outperform these models using data-driven and physics-informed approaches. ML control could potentially leverage various HET power supplies simultaneously to control the HET [22, 25]. Past research has also explored using ESNs for time series prediction of discharge current [17].

Active control of a HET is desired as it can shape oscillations to reduce wall erosion, reduce EMI, protect the PPU, improve stability, and improve thrust, specific impulse, and efficiency. [12–15]. We propose a ML control approach that can be used to make intelligent PPUs. This approach can allow HETs to operate optimally by allowing them to either reduce oscillations or phase synchronize voltage and current to increase thrust. The approach operates faster than the BM oscillation timescale and uses an NMPC controller on a NN that could potentially update over time. NMPC allows

for predictive rather than reactive control and provides flexibility for control objectives as it classically involves solving a nonlinear optimization problem given a dynamics model and constraints. With increasing data availability, NNs can better generalize the dynamics of the HET discharge plasma and aid in computing better control trajectories [26]. ML control could also potentially predict and prevent mode-hopping behavior in HETs, where the HET switches between lower-thrust plume modes and higher-thrust jet modes [27]. The biggest benefit from ML control could be a technique that could unlock higher power operating points that have higher thrust, specific impulse, and efficiency that are not feasible due to the large oscillations that are tied to them by reducing them.

III. Control Problem and Approach

We start by modeling the discharge circuit of the HET and identifying the control problem. We then discuss our active control approach by investigating time and frequency domain behavior that we seek to change. Lastly, we discuss three different ways to control a HET that we argue are best achieved using ML rather than PID. The work presented here is broken into control simulations on a 0D ionization model, experimentation to obtain a HET dataset including PID control, training, testing, and tuning of an ESN using the dataset, PID and NMPC control simulations using the ESN, and deployment onto an FPGA for use in an experiment to demonstrate the ML controller for oscillation reduction.

The circuit in Figure 1 consists of a DC voltage source (the PPU output), a RC low-pass discharge filter to protect the PPU from oscillations, a modulation voltage, a harness transmission line model (the source of discharge voltage oscillations via time-varying current flowing through LH), and the load represented by a variable current source [28]. The following control model in Equation 1 shows that there is coupling between discharge current, modulation voltage, and discharge voltage that requires modeling. For minimizing oscillations, the control objective is to find a trajectory where VM opposes the natural oscillations in VD. However, the choice of control voltage is not as simple as negating the discharge voltage, as the discharge voltage changes with the applied control voltage.

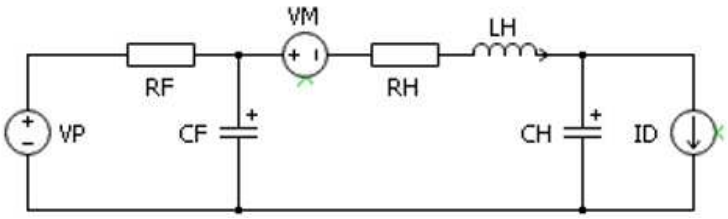


Fig. 1 Discharge Circuit Diagram

$$V_{CF} - L_H \frac{dI_D}{dt} - V_M - V_D = 0 \tag{1}$$

Figure 2 shows the FFT of the discharge current of the H9 9-kW HET operated at 600 V, 15 A. The FFT at this operating condition reveals a BM peak and harmonics. It is important to note that it is not the case that the breathing mode and harmonics are solely responsible for the oscillations; the sidebands of the peaks also are signals and contribute to the oscillatory behavior. Control does not involve just addressing the peaks, so it is necessary to develop a controller that can handle all oscillations depicted in the FFT. The rationale for control involves active noise cancellation, where waves at various frequencies are canceled by introducing waves that are 180 degrees out of phase, leading to destructive interference [29]. However, in a nonlinear load such as a HET, perturbations at a particular frequency can lead to responses at other frequencies due to frequency coupling, which need to be accounted for.

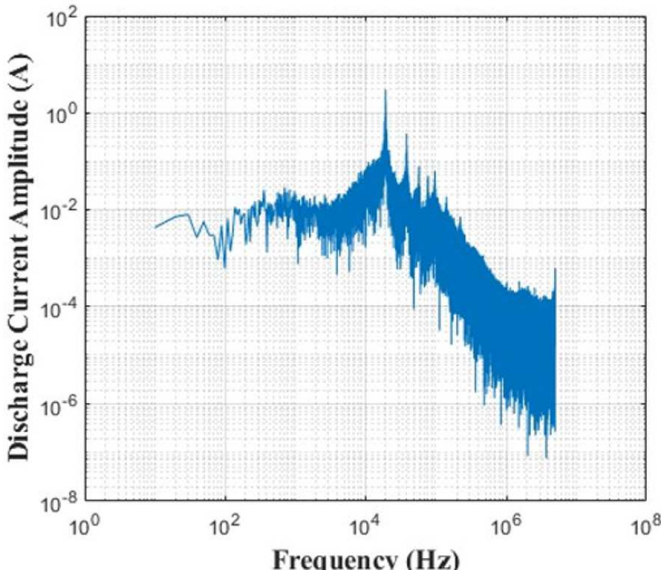


Fig. 2 Discharge Current FFT

Figure 3 shows the philosophy of active control, where the discharge current waveform is shaped through repeated applications of control voltages throughout the BM period. The proper choice of control voltages at the right points in time could allow for oscillation reduction. There are fundamental physical limits on the frequency of the control voltages, due to the behavior of ions, electrons, and neutrals.

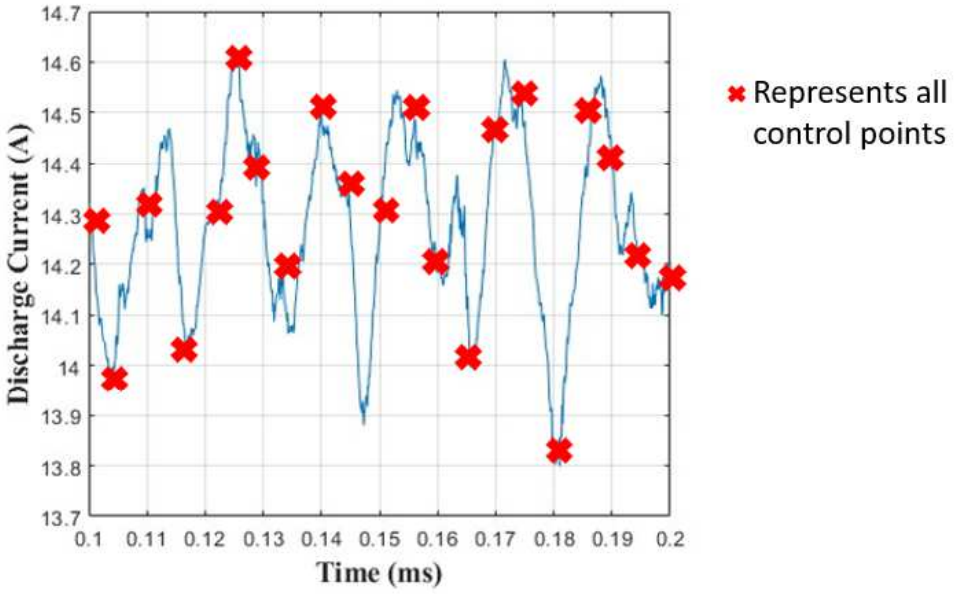


Fig. 3 Waveform Shaping of HET Discharge Current Oscillations

Table 2 shows three possible control modes that could be used for a HET. The first is to perturb the discharge voltage to control voltage and current oscillations to reduce wall erosion, reduce EMI, improve stability, and protect the PPU. The second is to perturb the discharge voltage to phase align the discharge voltage and discharge current to increase thrust. The third involves perturbing the discharge voltage to convert discharge current oscillations into AM/FM/PM, as the discharge plasma emits EM waves, which could potentially be used for communications.

Table 2 Table of Control Modes

Control Mode	Description
Voltage and Current Oscillation Control	Perturb voltage to control the voltage and current oscillations
Maximum Thrust	Achieved by phase aligning the voltage and current oscillations
AM/FM/PM	Achieved by applying appropriate voltage perturbations that result in AM/FM/PM of discharge current

From an electrical point of view, the approach extends selecting specific RLC components for passive power factor correction to active power factor correction by introducing voltage perturbations. This approach is a drive circuit for HETs. Figure 4 shows that the phase shift between discharge voltage and discharge current dictates the real and reactive power. For maximum thrust and real power, the voltage and current signals should be phase-aligned, minimizing reactive power and maximizing power factor. A 90-degree phase shift results in minimal real power input and, thus, minimal thrust. Rather than minimizing oscillations, it may even be advantageous to increase oscillations while synchronizing discharge voltage and discharge current, or even reduce oscillations while achieving synchronization. It is important to note that while the DC settings of the thruster largely impact thruster performance, the AC behavior also plays a role, and shaping AC behavior can lead to different outcomes in terms of performance.

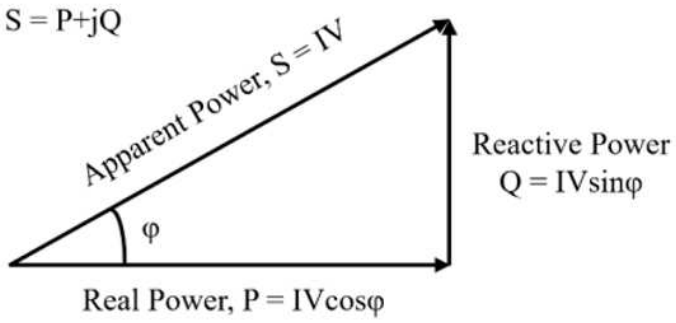


Fig. 4 Power Triangle

In this work, we choose to focus on oscillation reduction before looking into improving thrust or demonstrating AM/FM/PM. The nonlinear optimal control problem for reducing the RMS of the discharge current oscillations is shown as follows, where x_k represents HET states at time k, u_k represents the control voltage at time k, f(x_k, u_k) represents HET dynamics at HET states and control voltage at time k, g(x_k, u_k) represents the discharge current output at HET states and control voltage at time k, N is the number of steps looking forward, j represents particular times during the N steps looking forward, μ is the mean discharge current, and u_{min} and u_{max} are minimum and maximum bounds on the control voltage:

Dynamics: $x_{k+1} = f(x_k, u_k), \quad i_k = g(x_k, u_k)$

Objective: $$\min_{\{u_k,\ldots,u_{k+N-1}\}} \underbrace{\sum_{j=0}^{N-1} \left(i_{k+j} - \mu\right)^2}_{\text{minimize RMS of oscillations}}$$

where $\mu = \dfrac{1}{N} \displaystyle\sum_{j=0}^{N-1} i_{k+j}$

Subject to:

$x_{k+j+1} = f(x_{k+j}, u_{k+j}), \quad j = 0, \ldots, N-1,$

$u_{\min} \leq u_{k+j} \leq u_{\max}, \quad j = 0, \ldots, N-1,$

x_k given.

IV. Methods

Section A discusses how we perform system identification on a HET using voltage perturbations. Section B discusses the ML model we selected for representing the relationship between control voltage and discharge current based on the perturbation response data collected during system identification. Section C explains the experimental setup that we use for collecting perturbation response data and control. Section D explains the PID control experiment that we conducted on the HET load that also is used for creating the ML model based on perturbation responses from applying the PID control. Lastly, Section E discusses how we use the ML model as a part of an NMPC framework in a simulation where optimal control voltages are determined that minimize discharge current oscillations.

A. System Identification

To control the HET discharge plasma, we must first determine a model for HET discharge plasma dynamics through a process known as system identification, which we can then use to determine control trajectories that allow us to achieve our control objectives. We use an Agilent 33522A function generator as the signal source, an AE Techron 9110 switch-mode power amplifier, and two AE Techron T2000 transformers in parallel for isolation, each with primary to secondary turns ratio of two to one, to interface with the discharge circuit for load perturbation, shown in Figure 5. The power amplifier and transformer each have a bandwidth of 250 kHz. Each transformer is rated for 500 W, allowing up to 1 kW to be added to the plasma. The power amplifier is responsible for amplifying the voltage of the input signal while also increasing its current. It is important to note that the power amplifier does not provide current to the load; it must source the current needed to maintain the desired voltage across the transformer. The voltage across the load is what sets the current through the load, which is mirrored to the primary of the transformer from the secondary, setting the power amplifier current. The additional current is provided by the discharge filter capacitor. DC offset cancellation is used since DC will saturate the transformer, the difference amplifier is included to boost the signal to the highest allowable level of the power amplifier, and the bandpass filter is used to limit the injection signal frequency between the cutoff of the discharge filter and the power amplifier and transformer bandwidth. This setup effectively places a variable voltage source in series with the HET. The system can handle the plasma inrush current and fault conditions. PPU output voltage changes are limited by the discharge filter, so the injection path needs to be placed after the filter to enable perturbation and control at higher frequencies and fast time scales. However, for control at slow time scales, as in the case of the ILC approach, the PPU output voltage can be changed without attenuation from the discharge filter [22]. Control at faster time scales has added complexity but has the potential to be more effective.

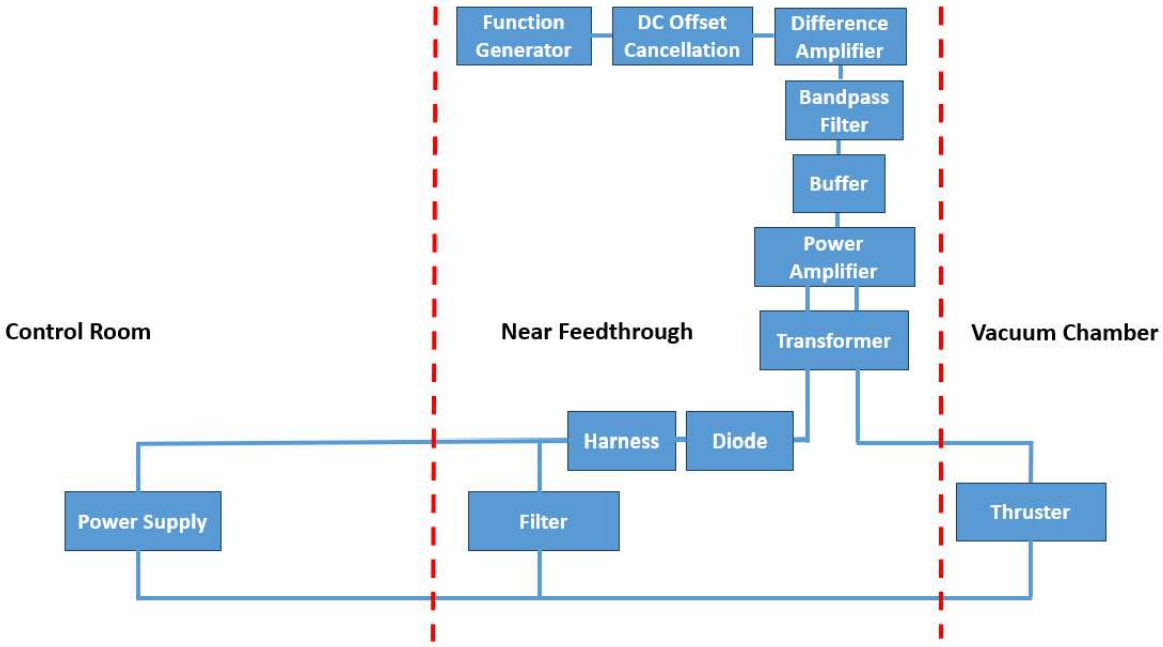


Fig. 5 Load Characterization Experiment Diagram

The perturbations we add to the discharge to assist with system identification of the discharge plasma are summarized in Table 3. These include sine waves for basic characterization of the plasma response, square waves to capture nonlinear responses, triangle waves which capture elements of both sine and square waves, ramp as a blend between triangle and square waves, chirps for characterizing time-varying effects, combinations of sinusoids to capture coupling, noise to capture stochastic effects, and lastly PRBS. These perturbations are applied at various amplitudes up to 2.5 V, frequencies between 7 kHz and 19 kHz, and phases at the 400 V operating point. This process can be applied to any HET operating point to learn the dynamics and create a ML controller. The system identification approach is still applicable for operating points that exhibit mode-hopping behavior.

Table 3 Load Characterization Test Matrix

Waveform Perturbation Type (Various Amplitudes, Frequencies, and Phases)
Sine
Triangle
Square
Ramp
Chirp
Discharge Voltage (Combination of Sinusoids)
Noise
PRBS
PID

B. ML Model

The main motivation for constructing an ML model is because HET discharge plasma models are not predictive enough to allow for control on their own. ML allows us to use a data-driven approach and additionally a physics-informed approach if desired to model HET discharge plasma dynamics. ML models are intended to predict the deterministic parts of the discharge plasma while attempting to predict the random contributions from plasma turbulence. In constructing a

ML model of the discharge plasma, it is necessary to measure the underlying states governing the plasma, specifically their response to perturbations. These include electrical states in the discharge circuit using the perturbation approach discussed previously along with plasma states. Plasma diagnostics are used to measure the plasma states shown in Table 4 below. However, it is difficult to measure the plasma states in real-time in the discharge. Non-intrusive plasma diagnostics would have to be used so that the plasma is not perturbed from an additional boundary condition.

Table 4 Time Scales of Various Plasma Diagnostics

Plasma Diagnostic	Invasive?	Plasma Parameter	Time Scale
Faraday Probe	Yes	U_i, N_i	Microseconds
Laser Induced Fluorescence	No	U_i, N_i	Not Real-Time
Triple Langmuir Probe	Yes	T_e, N_i	Microseconds
Laser Thomson Scattering	No	T_e, U_e, N_i	Not Real-Time
Optical Emission Spectroscopy	No	T_e, N_n, N_i	Not Real-Time

It is also difficult to use a 3D plasma discharge model governed by PDEs and take 3D measurements for plasma state estimation using an EKF due to the complexity and time involved in solving the PDE system and obtaining an adequate dataset. 3D control would also require placing electrodes in various spatial locations around the discharge channel or even actively biasing the thruster body and keeper as well [30]. Controlling these electrodes would pose a computational challenge. A more feasible approach would be to use a 0D model, discussed in the Appendix, and take time-varying and spatially-independent electrical measurements. Physics-informed ML could involve learning the residuals between measurements and predictions from the physics model, and the prediction of these residuals could then be used for control. These residuals would be the unmodeled terms in the 0D ionization model, primarily anomalous electron transport. Others have looked into symbolic regression for determining the physical representation of unmodeled terms not captured in physical models for the discharge plasma, but this procedure is very computationally and time-intensive, thus NNs may be the better option for control applications [31]. However, tuning the EKF can be cumbersome and requires that the plasma state estimate be correct in order to achieve a proper predictive model. Thus, we opt not to use state estimation and stick with just electrical measurements. Furthermore, we use only the mapping between the control voltage and the discharge current for simplicity, though the other electrical states, such as discharge voltage, harness inductor current, and discharge filter capacitor voltage may contribute additional information that may further improve prediction accuracy. A ML model was built using an ESN, which is a form of reservoir computing and a type of RNN that excels at predicting nonlinear and nonstationary time series [32, 33]. The ESN consists of an input layer, a reservoir layer that uses random, fixed recurrent connections that are never trained, and an output layer that is trained using linear regression, and is represented by Equations 2-7 as follows:

$$\hat{h}(n) = tanh(W_{in}x(n) + W_{rec}h(n-1) + W_{fb}y(n-1) + b_{rec}) \tag{2}$$

$$h(n) = (1-\alpha)x(n-1) + \alpha\hat{h}(n) \tag{3}$$

$$y(n) = W_{out}h(n) + b_{out} \tag{4}$$

Ridge regression through Tikhonov regularization can be used to determine W^{out}:

$$Y_{target} = W_{out}X \tag{5}$$

$$W_{out} = Y_{target}X^T(XX^T + \beta I)^{-1} \tag{6}$$

Prediction is accomplished using W_{out}^T:

$$Y_{pred} = XW_{out}^T \tag{7}$$

The mapping between control voltage and discharge current can be used to determine sequences of control voltages that can minimize discharge current oscillations, shown in Figure 6. Measurements of discharge current can be directly fed into the ESN to make control voltage predictions. We pick ESNs over LSTMs as they can be trained, tested, and tuned much faster on CPUs without the need for GPUs while maintaining comparable prediction accuracy. LSTMs have the potential to be better at learning temporal dependencies, but are more complex than ESNs [32]. ESN stability can also be controlled more easily than MLPs and LSTMs [34].



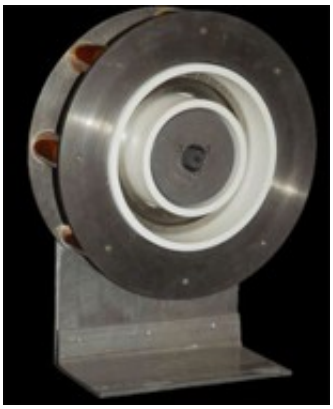
Fig. 6 ESN Model of Control Voltage and Discharge Current

Given the diverse set of perturbations applied to create the training dataset, the ESN can better learn signal from noise since the noise is smoothened. A lock-in amplifier could also potentially be used to improve the SNR of the measurements. Due to the presence of noise, the true discharge current measurements actually look different from what one may observe on an oscilloscope. It is important to control signal rather than noise, and the effectiveness of the control is reduced when more noise is present. Higher power operating points at higher voltages and currents tend to have higher amplitude components with higher SNR, which further improves control performance, also taking into account that higher power generally leads to more erosion and larger oscillations that can be synchronized for increased thrust [15].

C. Experimental

Experimentation to obtain a perturbation dataset for developing ML control was performed in VTF-2 at the Georgia Institute of Technology High-Power Electric Propulsion Laboratory. An photograph of this facility is shown in Figure 7b. VTF-2 is a stainless-steel chamber measuring 9.2 m in length and 4.9 m in diameter. VTF-2 is evacuated to a rough vacuum using a 495 CFM rotary-vane pump and a 3800 CFM blower. High vacuum is achieved using ten liquid nitrogen-cooled CVI TMI re-entrant cryopumps at a pumping speed of 350,000 l/s on xenon. The cryopump shrouds are fed using the Stirling Cryogenics SPC-8 RL special closed-loop nitrogen liquefaction system [35].

For this study, we used the H6 HET, shown in Figure 7a, a 6-kW class HET developed in partnership between the NASA Jet Propulsion Laboratory, the University of Michigan, and the Air Force Research Laboratory [1]. The H6 employs a centrally mounted LaB_6 cathode. The thruster body and cathode were isolated from facility ground and were electrically floating.



(a) H6 HET.



(b) VTF-2.

Fig. 7 Test Article and Facility

The H6 discharge plasma was controlled using a Magna-Power TS800-24 power supply. All other thruster components were powered using TDK-Lambda GEN80-42 power supplies. A TDK-Lambda GEN600-2.6 and GEN60-25 were used for the cathode keeper and heater, respectively. To protect against HET oscillations, the discharge supply was connected to a RC low-pass discharge filter consisting of a series 0.533-Ohm resistor and a shunt 100-µF capacitor to attenuate discharge current oscillations that are greater than 3 kHz in frequency.

High-purity (99.999%) krypton propellant was supplied to the anode and cathode lines using stainless-steel lines metered with MKS GE50A MFCs. The MFCs were calibrated by measuring the flow upstream of the thruster with

a MesaLabs DryCal 800-10 volumetric flow meter. The uncertainty for anode and cathode line flow is 2% and 5%, respectively.

The power supply voltage, anode and cathode flow rate, and the magnetic field were kept constant during testing. The H6 thruster was operated at 400 V, 4.3 A and 250 V, 4.3 A thruster operating points, as shown in Table 5. Upon ignition, the H6 thruster was run for three hours until its oscillations became steady after outgassing. The base facility pressure was 10 nTorr, and the facility pressure during operation was 2 µTorr.

Table 5 H6 Thruster Operating Conditions (Krypton)

Discharge Voltage (V)	Discharge Current (A)	Anode Flow Rate (sccm, mg/s)	Cathode Flow Rate (sccm, mg/s)	Inner Coil Current (A)	Outer Coil Current (A)	Breathing Mode Frequency (kHz)
400	4.3	65, 4.05	6.5, 0.405	2.15	1.96	13
250	4.3	65, 4.05	6.5, 0.405	2.15	1.96	6.646

The AC component of the discharge current of the H6 thruster was recorded using a Pearson 110A current monitor placed inside VTF-2, and the control voltage was measured using a BNC cable connected to a HDO6104 Teledyne LeCroy oscilloscope. The uncertainty and bandwidth of the current probe are ±1% and 20 MHz; for the oscilloscope, they are ±0.5% full scale and 1 GHz. The oscilloscope has 12 bits of resolution with a selected 25 MS/s sampling rate and is used over an interval of 10 ms for 250,000 samples in all. MATLAB instrument control toolbox with USB and Ethernet for remote desktop of the oscilloscopes were used for semi-automated data collection. HET measurements were collected after it was observed that the peak-to-peak current values had stopped changing within ±1% variation, taking approximately two hours after ignition due to outgassing.

Systematic and random errors were quantified as part of this study's error analysis. Systematic errors are introduced by voltage and current probes, along with the oscilloscope. Other errors and nonidealities arise from function generator, microcontroller, ADC, FPGA, DAC, power amplifier, transformer, and harness settings.

D. PID Control Experiment

As a part of determining perturbation responses, we conducted PID control of the HET discharge plasma. The premise behind PID is that the discharge current oscillations can be reduced by minimizing the error between current measurements and a reference. This is achieved by supplying the error into the PID network, which computes the control voltages needed to accomplish the minimization of the error, represented in Equation 8.

$$V_{AC} = K_P I_{AC} + K_I \int_0^t I_{AC}\, dt + K_D \frac{dI_{AC}}{dt} \tag{8}$$

The control scheme works by applying voltage at 10 microsecond steps while taking into account propagation and computation delays, shown in Figure 8. For the control approach to be effective, propagation delays between the current probe and the DAC along with computation delays must be minimized.

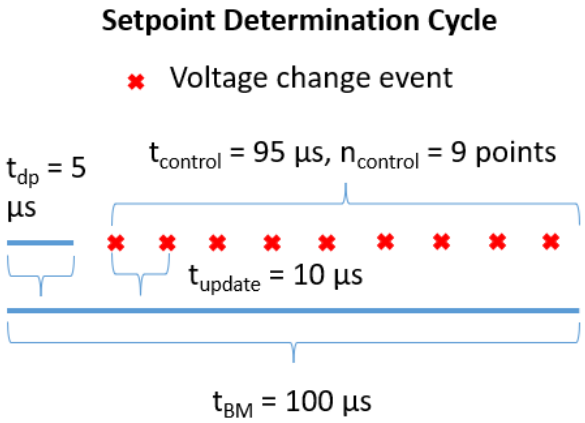


Fig. 8 Setpoint Timing

PID control was achieved using the Texas Instruments C2000 F28379D microcontroller. The C2000 embedded coder target in Simulink was used, allowing for rapid programming of the microcontroller using various blocks that are used to control it, in addition to the standard PID block. The embedded coder provides functionality to set the control loop speed. The microcontroller has a built-in ADC to sample the discharge current and a DAC to output the appropriate control voltage after PID computation, ranging between 0 and 3.3 V. A BNC switcher is used to protect the microcontroller from the initial inrush voltage signal from the Pearson current probe used to measure the discharge current of the HET. Resistive dividers are used to set appropriate voltages, and operational amplifiers are used for buffering, scaling, and level shifting the measurements to fit within the input range of the ADC.

Figure 9 shows the experimental setup for PID control, where the Texas Instruments C2000 microcontroller replaces the function generator from the system identification experiment. We treat all components after the DAC combined to be the plant. The plant notably does not include the RC discharge filter, which contains a 50 μs time delay, since the propagation delay from the transformer through thruster harnessing to the HET is what is relevant in this setup. Nonetheless, the controller computes the control voltages that take into account any delays so that the discharge current can be controlled as desired.

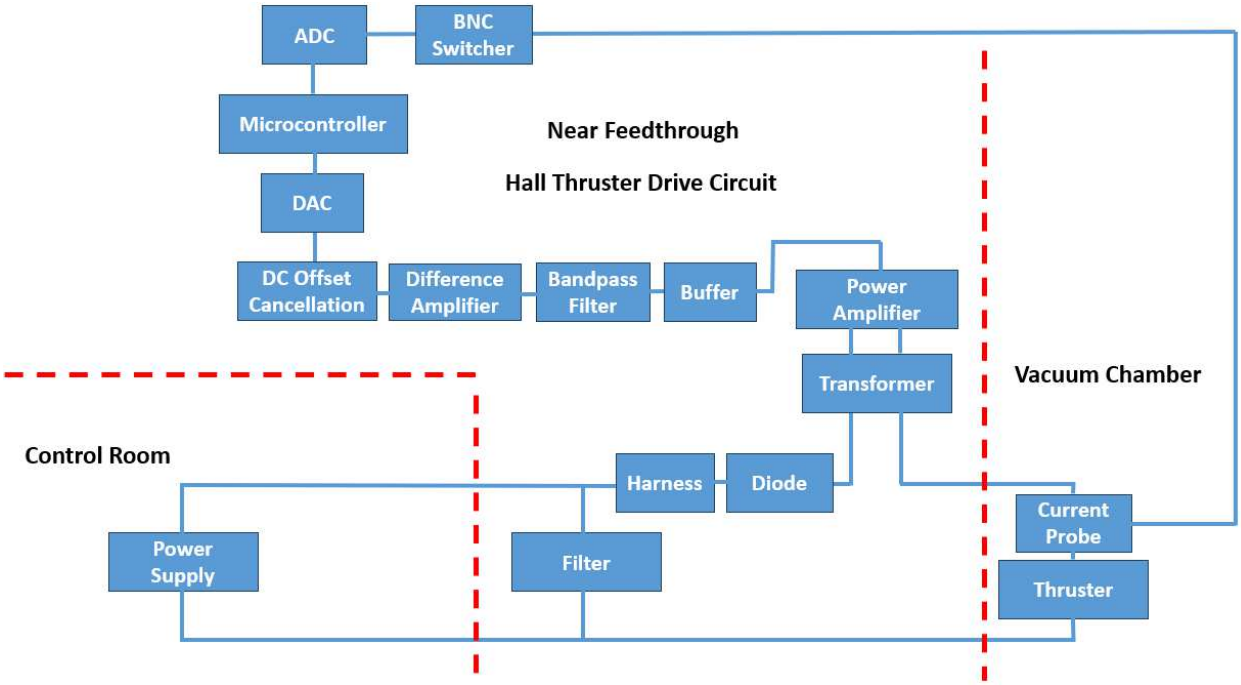


Fig. 9 Experimental setup for PID Control of the HET load

E. ESN NMPC Control Simulation

The ML controller we seek to develop consists of an ESN and an NMPC controller [33]. Figure 10 shows the basic architecture of the ML controller, where discharge current values are read and a corresponding control voltage is identified to limit the discharge current oscillations. The ML model in the controller can also be updated over time to close the loop. To solve the nonlinear optimal control problem using NMPC, we use the IPOPT nonlinear optimization approach to determine optimal control trajectories over a horizon that minimize discharge current oscillations [36]. We must use a nonlinear solver for the ESN since the ESN is nonlinear. IPOPT uses an interior point method with KKT, exploits sparsity, and can be coupled with CasADi for exact derivatives [37]. IPOPT is slower than acados, another nonlinear optimization solver, but is more accurate as it better handles stiff dynamics and constraints. IPOPT is also a better choice than SQP, since SQP also struggles with stiff dynamics.

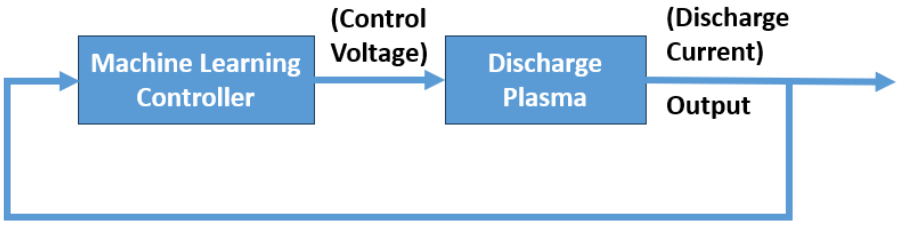


Fig. 10 ML Control Feedback Diagram

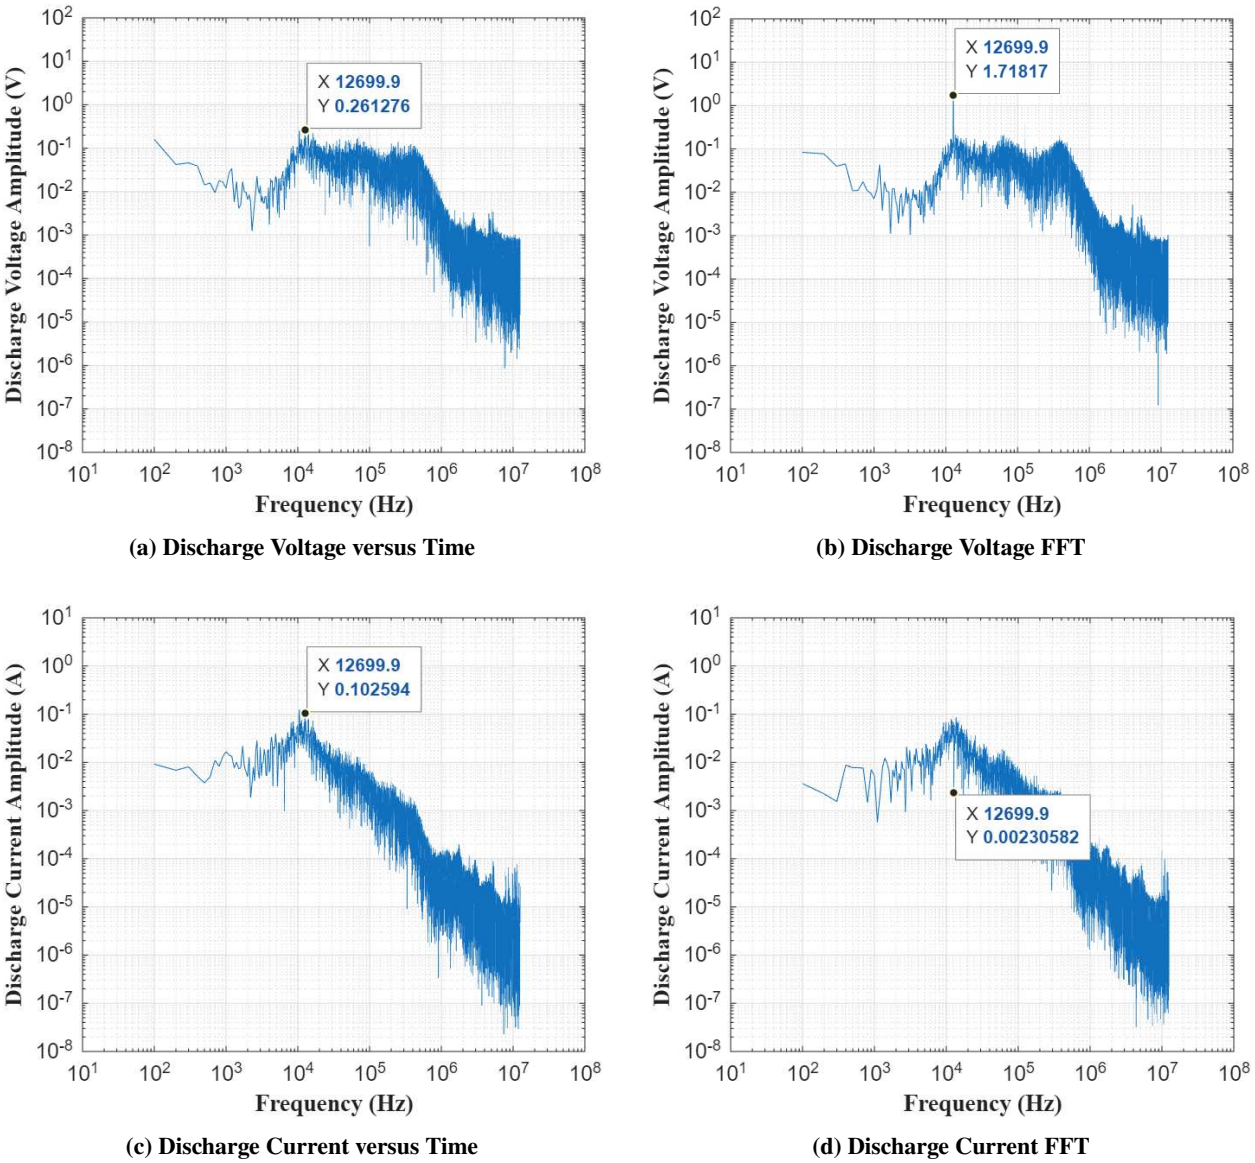
V. Results and Discussion

In our work, we seek to compare the performance of NMPC-based ML control approach with PID. We present results and discuss oscillation reduction achieved using PID control versus without using PID control. We show that the perturbations actually cause changes in the discharge current using the eCCM algorithm. We show plasma state estimate oscillations using an extended Kalman filter and discuss the resulting plasma physics from PID control beyond electrical changes. We then show the results from ESN training and testing that demonstrate that ESNs can represent discharge current dynamics of a HET. We lastly show that PID and NMPC can control the discharge current oscillations in simulation using the ESN.

A. Load Characterization Experiment Results

1. 400 V Operating Condition

The perturbations at this operating point did not produce consistent observable response, with amplitudes comparable to noise. However, we did observe destructive interference with a 2 V, 12.7 kHz sinusoidal perturbation. Figure 11a shows the voltage FFT without perturbation, Figure 11b shows the voltage FFT with perturbation, Figure 11c shows the current FFT without perturbation, and Figure 11d shows the current FFT with perturbation. Figure 11d clearly shows a much lower amplitude compared to Figure 11c. Destructive interference, along with constructive interference, do not always happen in a plasma and requires proper timing.



(a) Discharge Voltage versus Time

(b) Discharge Voltage FFT

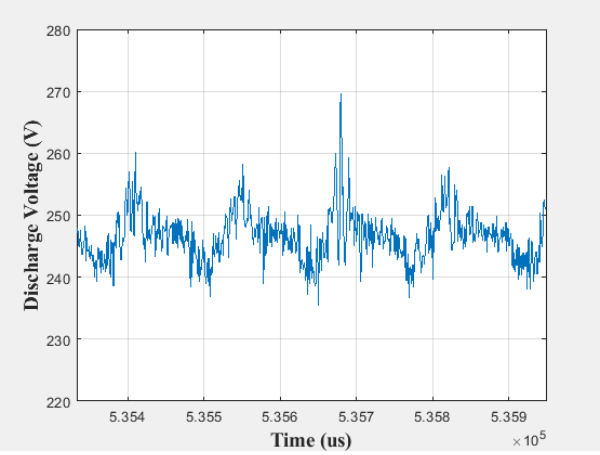
(c) Discharge Current versus Time

(d) Discharge Current FFT

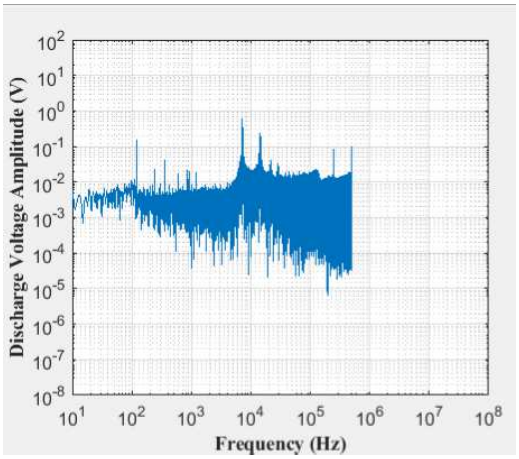
Fig. 11 Discharge Voltage and Discharge Current at 400 V Operating Point With 2 V Sinusoidal Perturbation at 12.7 kHz

2. 250 V Operating Condition

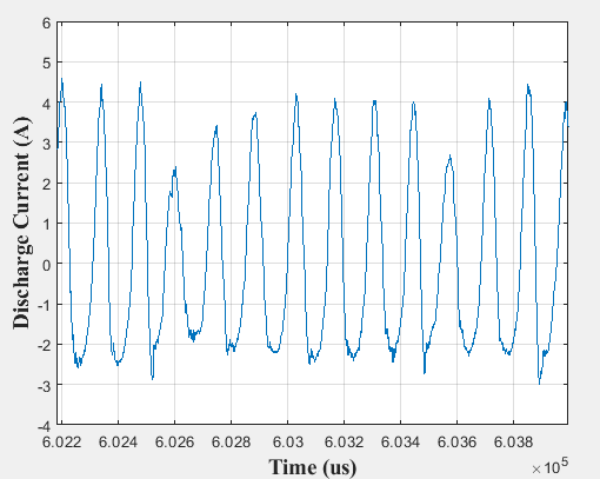
Figures 12a and 12b show the discharge voltage time series and FFT with no voltage perturbation. Figure 12a shows a triangle wave-like waveform in the time domain with a peak-to-peak voltage of roughly 10 V, with noticeable overshoots and ringing. Figure 12b shows the FFT, which shows a breathing mode fundamental at 6.646 kHz, along with several harmonics. Figures 12c and 12d show the corresponding discharge current times series, with a peak-to-peak of around 6 A, along with its FFT.



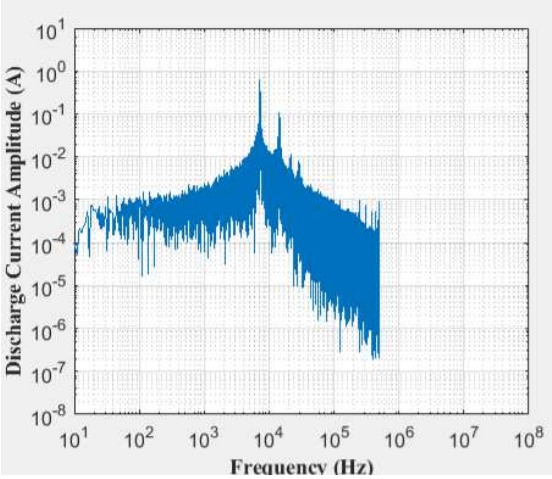
(a) Discharge Voltage versus Time



(b) Discharge Voltage FFT



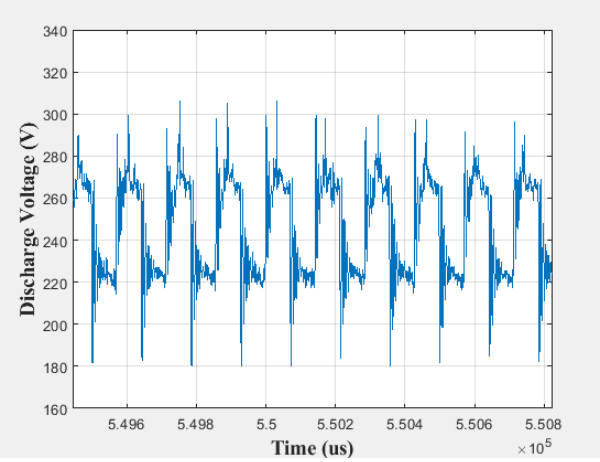
(c) Discharge Current versus Time



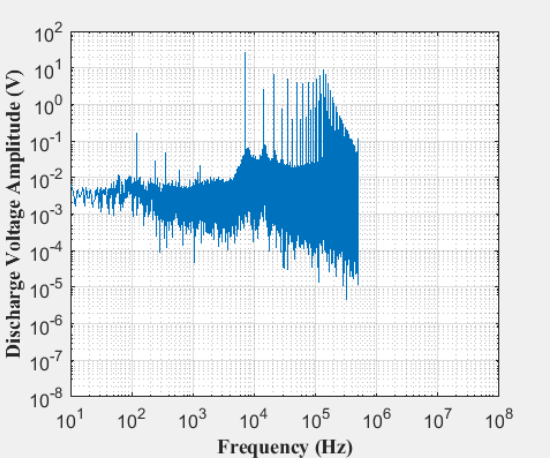
(d) Discharge Current FFT

Fig. 12 Discharge Voltage and Discharge Current at 250 V Operating Point

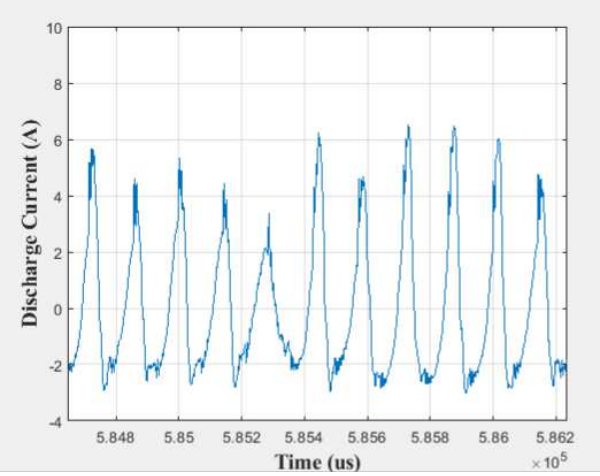
Figures 13a and 13b show the discharge voltage time series and FFT in response to a 20 V, 7 kHz, 0 degree phase square wave voltage perturbation. As opposed to the unperturbed discharge voltage time series observed in Figure 12a, the response in Figure 13a shows a square wave-like waveform in the time domain with a peak-to-peak voltage of roughly 50 V, with noticeable overshoots and ringing. Figure 13b shows the FFT, which shows many square wave harmonics that were not seen in Figure 12b. Figures 13c and 13d show the corresponding discharge current response, with a peak-to-peak of around 8 A. This result again shows that it is possible to perturb a plasma at many frequencies and observe responses.



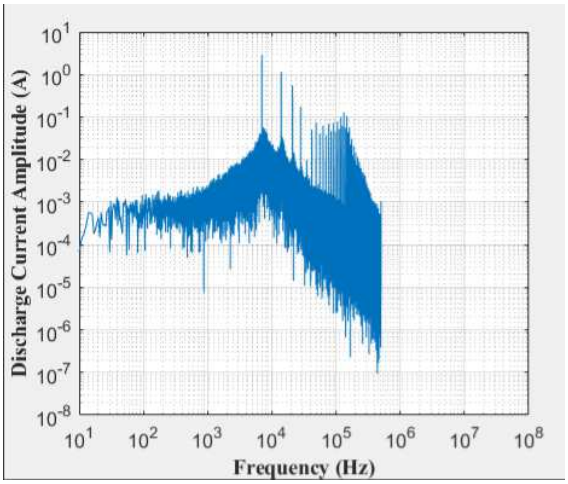
(a) Discharge Voltage versus Time Response



(b) Discharge Voltage FFT Response



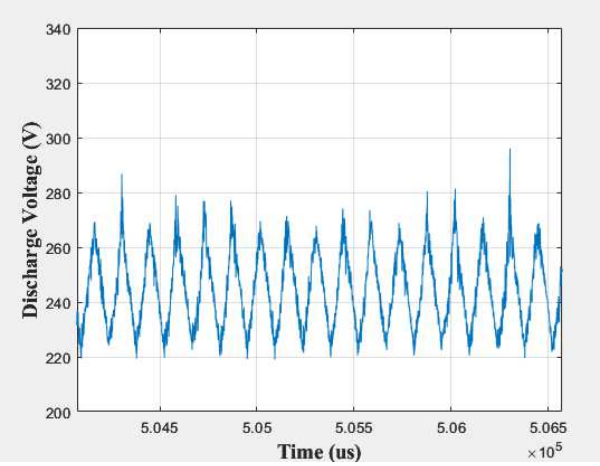
(c) Discharge Current versus Time Response



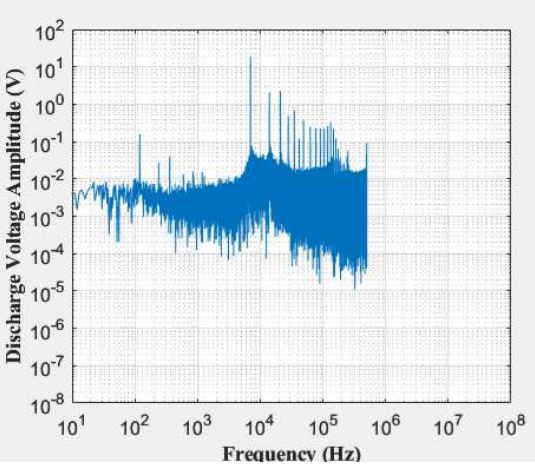
(d) Discharge Current FFT Response

Fig. 13 Square Wave Responses at 250 V Operating Point

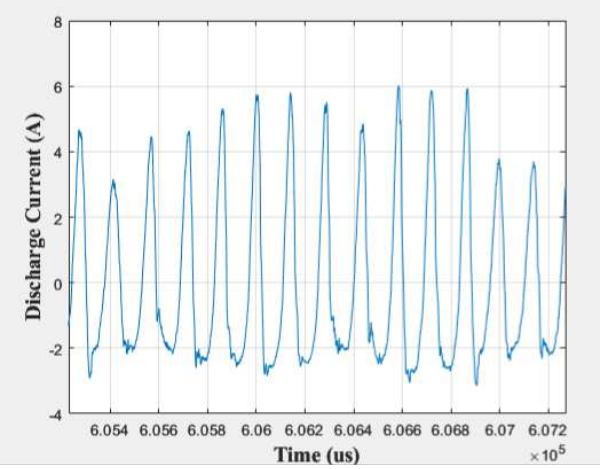
Figures 14a and 14b show the discharge voltage time series and FFT in response to a 20 V, 7 kHz, 0 degree phase triangle wave voltage perturbation. As opposed to the unperturbed discharge voltage time series observed in Figure 12a, the response in Figure 14a shows a largely triangular waveform in the time domain with a peak-to-peak voltage of roughly 60 V, with some noticeable overshoots. Figure 14b shows the FFT, which shows many triangle wave harmonics that were not seen in Figure 12b. Figures 14c and 14d show the corresponding discharge current response, with a peak-to-peak of around 8 A. This result again shows that it is possible to perturb a plasma at many frequencies and observe responses.



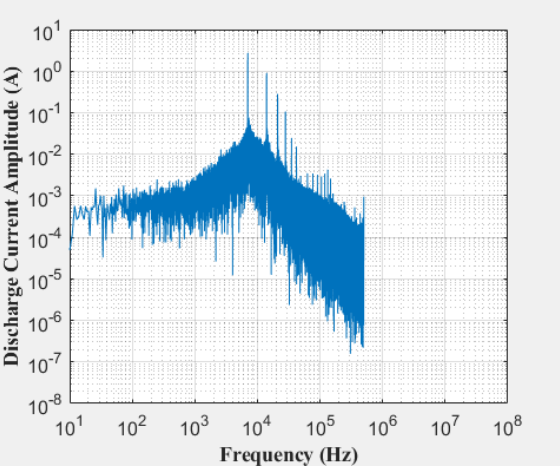
(a) Discharge Voltage versus Time Response



(b) Discharge Voltage FFT Response



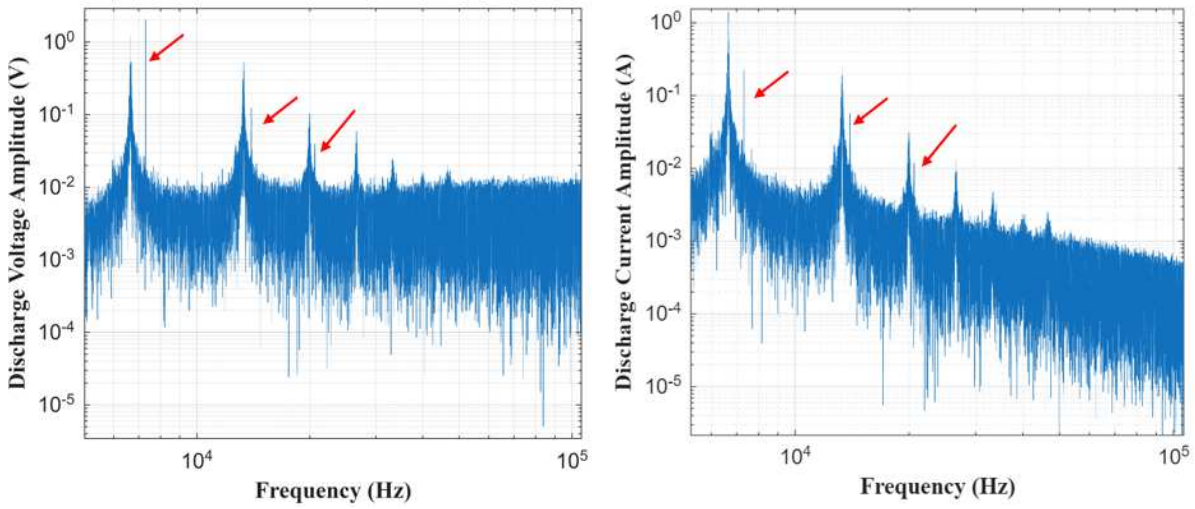
(c) Discharge Current versus Time Response



(d) Discharge Current FFT Response

Fig. 14 Triangle Wave Responses at 250 V Operating Point

Figure 15a shows the discharge voltage FFT in response to a 5 V, 7 kHz, and 0 degree phase sinusoidal voltage perturbation, with a visible vertical peak at this frequency. It can be observed that in addition to the prexisting breathing mode fundamental and harmonics, the perturbation generates harmonics of its own, positioned to the right of each breathing mode harmonic peak. The same behavior is mirrored in the discharge current FFT. This finding has major implications for HET control, as only control of the breathing mode could be necessary, resulting in control of the breathing mode harmonics as well. The same behavior exists for perturbations on the left side of the breathing mode peak. This behavior stems from the nonlinear nature of the breathing mode, where one frequency input creates additional frequencies.



(a) BM and Harmonic Voltage Response **(b) BM and Harmonic Current Response**

Fig. 15 BM and Harmonic Response at 250 V Operating Point

Figure 16 shows a bandlimited gaussian signal as seen at the function generator output. The voltage to be set across the Hall thruster is 25 V, however filters and other drops along the way to the load reduce the actual voltage seen at the load. The bandlimited gaussian is a signal that contains frequency content centered around the breathing mode, ranging from 6 kHz to 7.3 kHz. The signal has a gaussian distribution, and thus includes a rich set of instantaneous phases, both negative and positive. This signal was constructed with the intent of perturbing the breathing mode with various phases so that constructive and destructive interference could be observed and used for training a NN model of the discharge plasma.

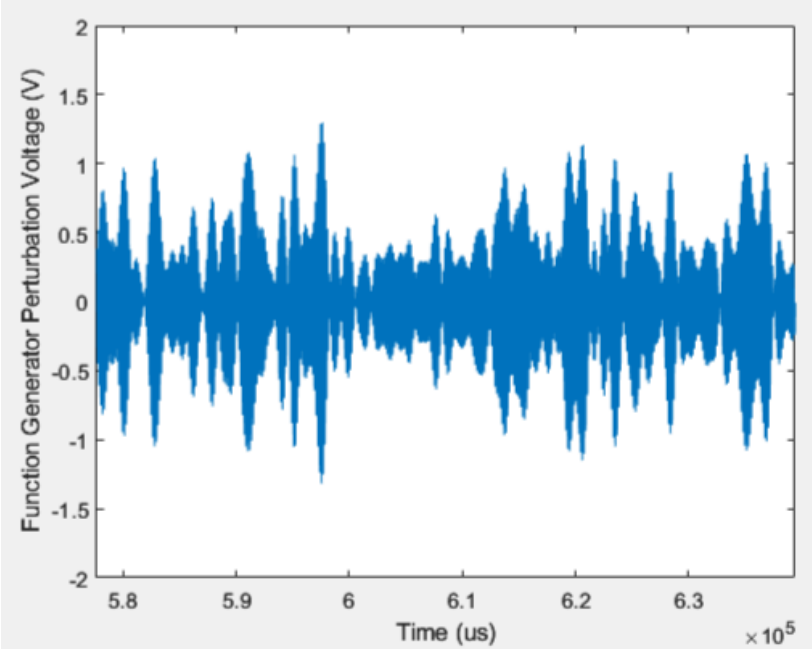
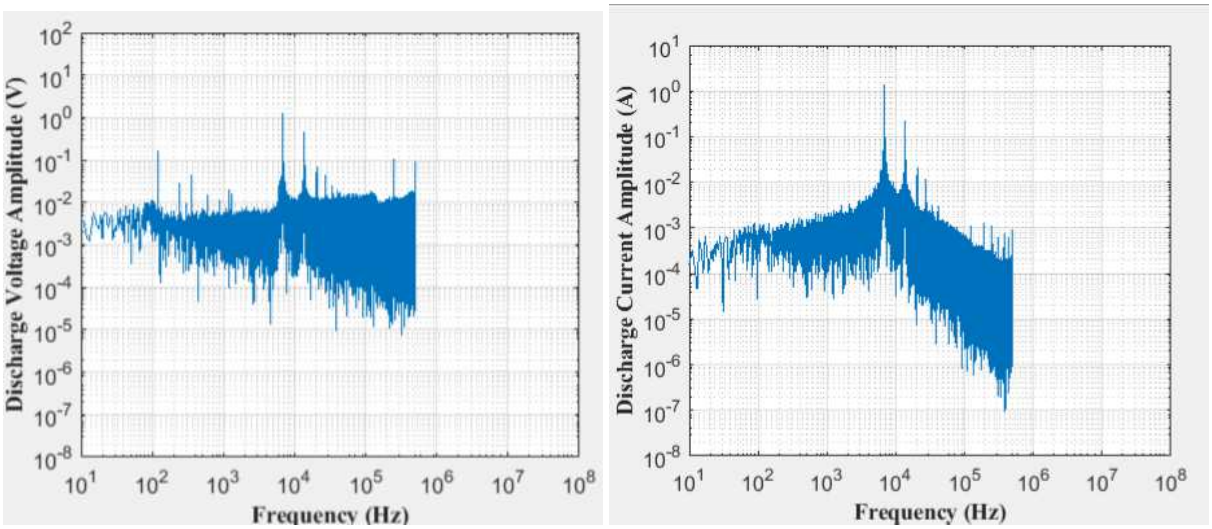


Fig. 16 25 V Bandlimited Gaussian at 250 V Operating Point

Figure 17a shows the discharge voltage FFT in response to the 25 V bandlimited gaussian voltage perturbation. Figure 17b shows the discharge current FFT in response to the perturbation. At this level of perturbation, there does not

appear to be a significant difference from the unperturbed response.



(a) Discharge Voltage FFT Response **(b) Discharge Current FFT Response**

Fig. 17 25 V Bandlimited Gaussian Responses at 250 V Operating Point

Figure 19a shows the discharge voltage FFT in response to the 150 V bandlimited gaussian voltage perturbation, shown in 18. Again, the intended voltage perturbation at the load is 150 V, though the actual voltage perturbation is less than 150 V. Figure 19b shows the discharge current FFT in response to the perturbation. It is clear that the voltage perturbation has changed the breathing mode voltage structure, and has reduced the breathing mode discharge current peak from over 1 A to 0.5 A, roughly a 50 % reduction. Again, we can see the same response mirrored in the breathing mode harmonics. This result shows great promise, demonstrating that bandlimited gaussian perturbation signals are excellent candidates for HET system identification.

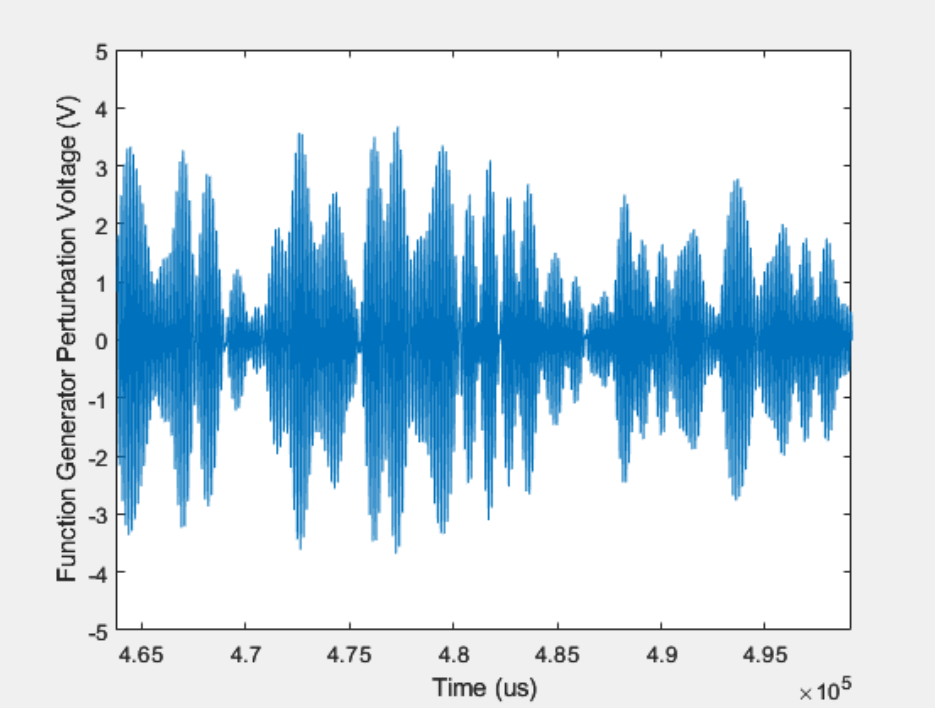
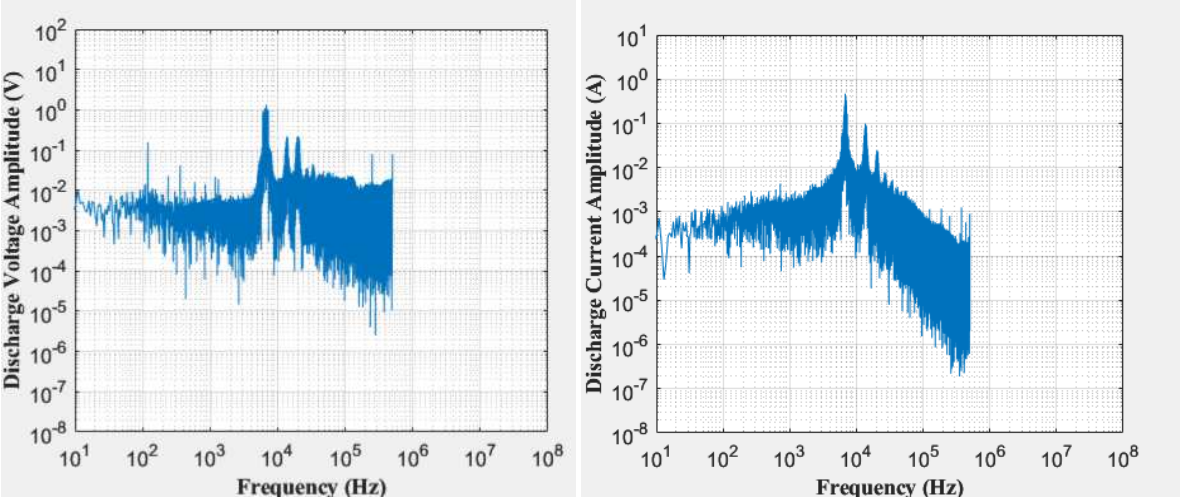


Fig. 18 150 V Bandlimited Gaussian at 250 V Operating Point



(a) Discharge Voltage FFT Response **(b) Discharge Current FFT Response**

Fig. 19 150 V Bandlimited Gaussian Responses at 250 V Operating Point

Figure 20 shows a 3D plot of the various sinusoidal perturbations that were applied to the HET discharge plasma. They span all combinations of the following sets of five amplitudes, five frequencies, and four phases, shown in Table 6, totaling 100 distinct perturbations that were used to develop a NN discharge plasma model. The frequencies include the breathing mode frequency of 6.646 kHz, and two frequencies to the left and right of the breathing mode peak:

Table 6 Voltage Perturbation Set at 250 V Operating Point

Frequency (kHz)	Amplitude (V)	Phase (Degrees)
6, 6.25, 6.646, 7, 7.3	1, 1.5, 2, 2.5, 5	0, 90, 180, -90

A second set of perturbations of eight amplitudes, three frequencies, and four phases taken on a day different from when the perturbation set was taken is also included in the plot and shown in Table 7, totaling 24 distinct perturbations. The frequencies include the breathing mode peak at 7.243 kHz and a frequency to the left and right of the peak. Because the data was taken during a different thruster firing with slightly different characteristics, we opted to not use this dataset for training the NN model of the HET load:

Table 7 Impedance Voltage Perturbation Set at 250 V Operating Point

Frequency (kHz)	Amplitude (V)	Phase (Degrees)
5, 7.243, 10	1, 1.5, 1.75, 2, 2.25, 5, 10, 25	0

It is important to note that this perturbation set is just a small fraction of the total perturbation space that is enclosed. For better ML control, it would be necessary to sample more perturbations while actively performing online learning from control output.

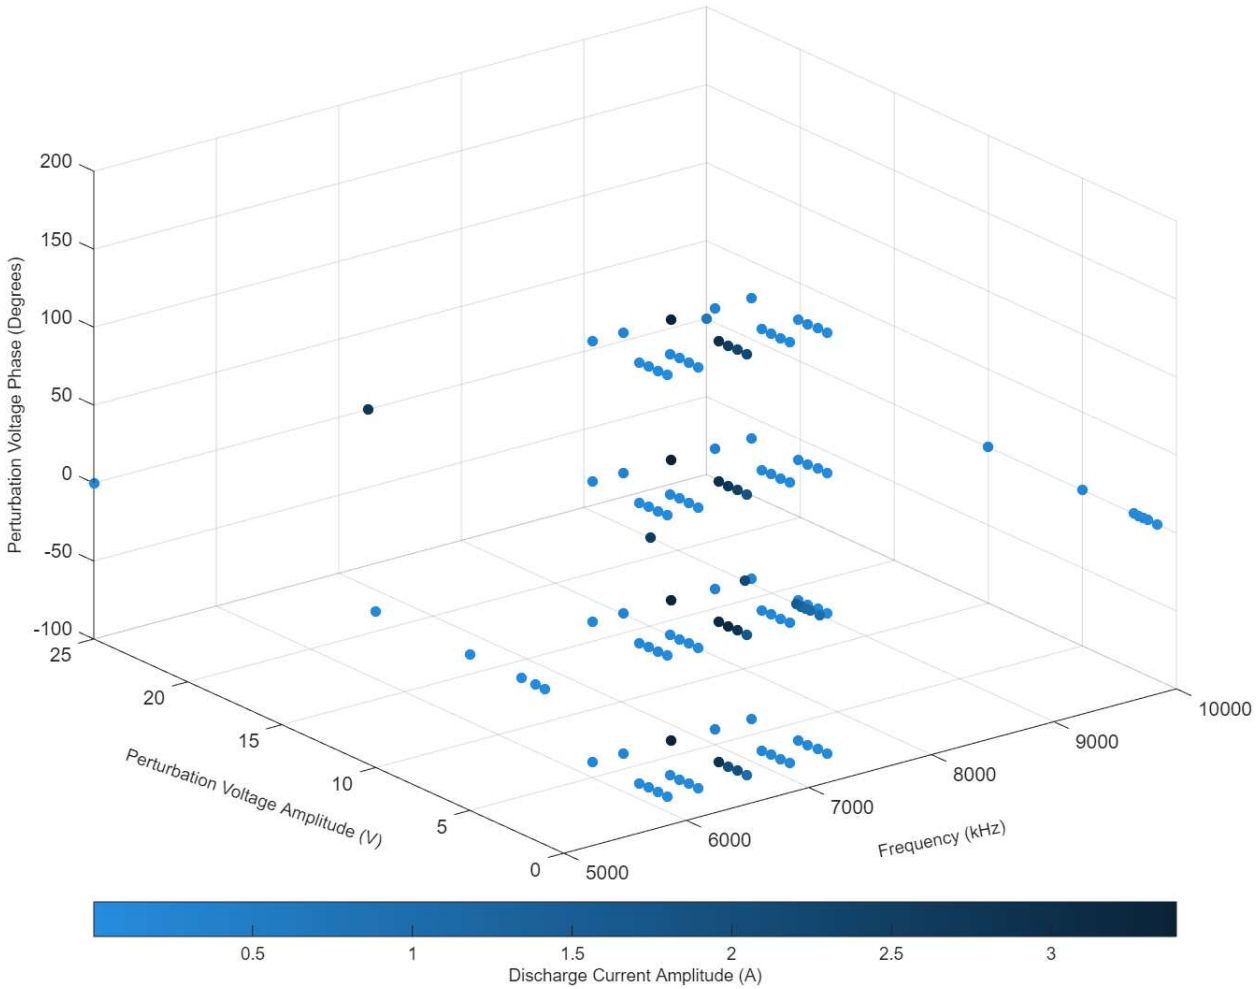
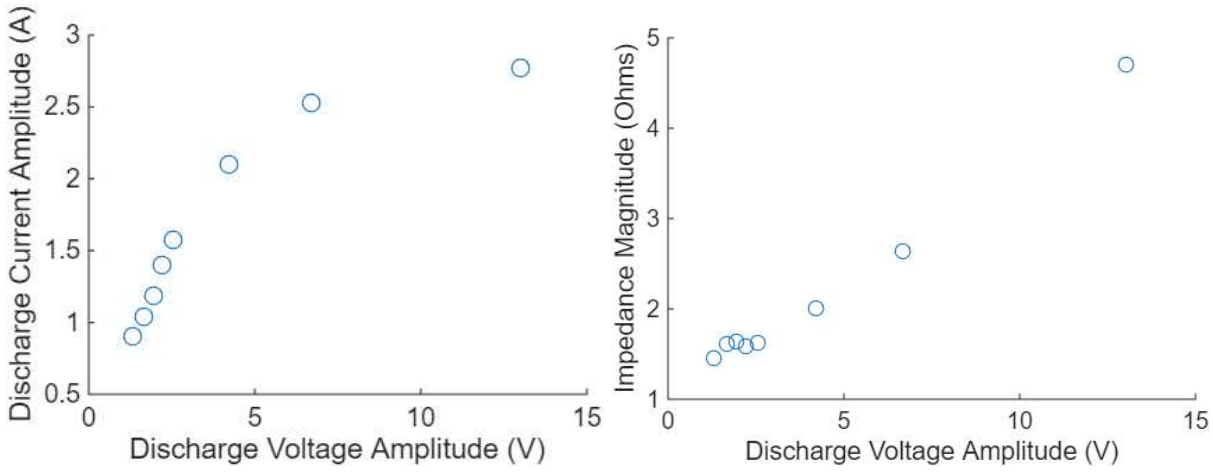


Fig. 20 Perturbation Voltage versus Current Map at 250 V Operating Point

Figure 21a shows the IV curve with sinusoidal perturbations at the breathing mode frequency at eight different amplitudes and a phase of zero degrees. We can see that the initial discharge current response at the same frequency is linear up to 2.5 V of discharge voltage, which is the limit of the small-signal linear regime of the plasma, as the curve then becomes nonlinear for discharge voltages above 2.5 V. It is interesting to note that the curve looks almost like a square root function, where the discharge current response begins to saturate as the discharge voltage increases. From a plasma physics standpoint, this makes sense as there is only so much current that can be drawn from ionization of the propellant. Figure 21b shows an alternate representation of the plot, plotting impedance against discharge voltage. We can see that the impedance is roughly constant for voltages below 2.5 V, but then grows, limiting the discharge current response. In terms of control, this result suggests that it is highly unlikely that large accidental perturbations, such as from momentary error from a NN, will damage the HET, and that control could be safe for use in space applications.



(a) Voltage versus Current Curve at BM **(b) Voltage versus Impedance Curve at BM**

Fig. 21 IV and Impedance Curves at 250 V Operating Point

B. PID Experiment Results at 400 V Operating Point

Table 8 shows the K_p and K_d values that were used during PID control of the HET discharge plasma. The integral coefficient was not used to keep the control simple and prevent integral windup. It can be observed that increasing the K_p value improves oscillation reduction, whereas increasing the K_d value makes the discharge voltage and discharge current more oscillatory. A K_p value of 1×10^{-1} and a K_d value of 1×10^{-9} minimizes the discharge current oscillations at 0.578 Arms, a 8.89% reduction, while also reducing the discharge voltage oscillations by 1.19% to 6.819 V_{RMS}. At worst, a K_p value of 5×10^{-2} and a K_d value of 5×10^{-3} maximizes the discharge current oscillations at 0.745 A_{RMS}, a 17.42% increases, while also increasing the discharge voltage oscillations by 58.02% to 10.908 V_{RMS}.

These results are consistent with the work done by Barral, where he demonstrated that the discharge voltage plays a significant role in the effectiveness of PID control [12]. Particularly, he found that increasing the discharge voltage mostly decreases the oscillation reduction percentage, likely due to the increasing nonlinearity of the HET at higher discharge voltages. His results show that operating conditions can play a role in the effectiveness of PID control, which we seek to overcome using ML control.

Table 8 PID Characterization at 400 V Operating Point

K_p	K_d	Discharge Current Oscillation RMS (No Control) (A)	Discharge Current Oscillation RMS (PID Control) (A)	Discharge Current RMS Change (%)	Discharge Voltage Oscillation RMS (No Control) (V)	Discharge Voltage Oscillation RMS (PID Control) (V)	Discharge Voltage RMS Change (%)
5×10^{-2}	1×10^{-9}	0.635 ± 0.034	0.635 ± 0.064	0.5 ± 11.4	6.902 ± 0.193	6.676 ± 0.059	-3.24 ± 2.84
5×10^{-2}	1×10^{-6}	0.635 ± 0.034	0.622 ± 0.031	-1.81 ± 5.8	6.902 ± 0.193	6.695 ± 0.063	-2.94 ± 3.05
5×10^{-2}	5×10^{-3}	0.635 ± 0.034	0.743 ± 0.07	17.42 ± 1.66	6.902 ± 0.193	10.9047 ± 0.292	58.02 ± 4.37
1×10^{-1}	1×10^{-9}	0.635 ± 0.034	0.5783 ± 0.03	-8.89 ± 5.45	6.902 ± 0.193	6.819 ± 0.142	-1.19 ± 2.21

Figure 22 shows the control voltages determined by the PID controller based on discharge current measurements at a K_p value of 1×10^{-1} and a K_d value of 1×10^{-9}, and hereafter we discuss PID performance at these PID gains. It shows a series of different levels of voltage at each 10 μs period, representing the 100 kHz loop speed. It is evident that the controller is making control voltage choices during each period of the breathing mode.

Due to the two parallel transformers used in the voltage perturbation approach, the control current required is the same as the discharge current. Reducing the discharge current oscillations reduces the required control current. We can compute the input power required to achieve a discharge current oscillation reduction of approximately 9% by taking the product of the RMS of the injected control voltage and the RMS of the discharge current. By scaling the control voltage by the gain introduced by the perturbation pathway, the RMS voltage at the output of the power amplifier is 12.221 V, while the RMS control current is 0.578 A, yielding a RMS control power of 7.064 W. Out of the 1.72 kW of DC power supplied to the HET, only 0.411% is needed to reduce discharge current oscillations by 9%, through proper selection of perturbation amplitudes and phases at each frequency.

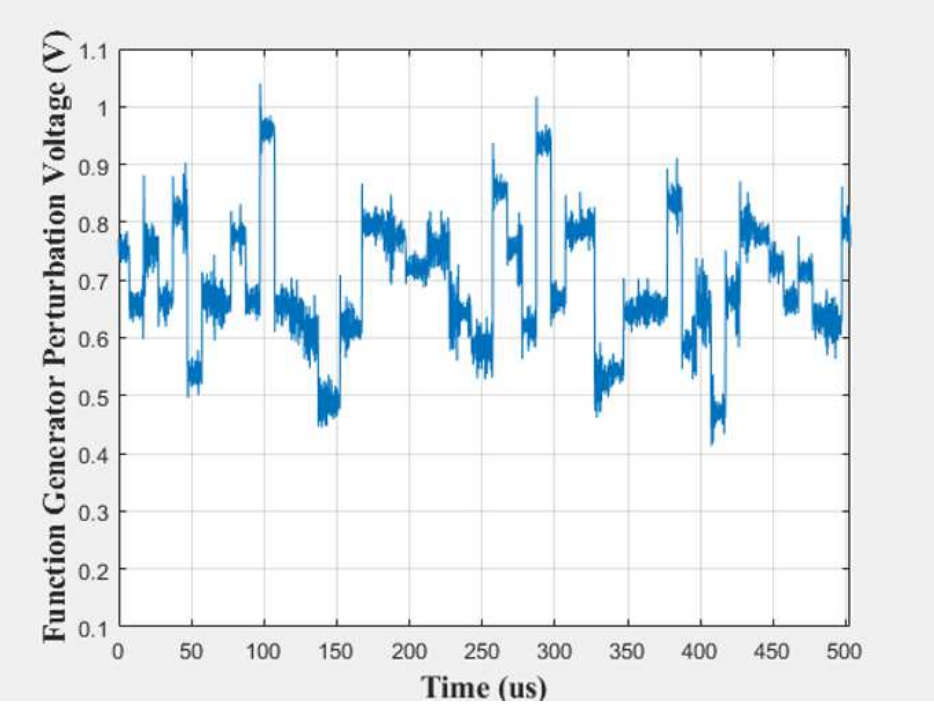


Fig. 22 PID Control Voltage at 400 V Operating Point

Determining correlation and causality between two signals in a nonlinear and time-varying system, such as a HET, can be accomplished using eCCM, a technique based on Takens' Theorem, and using shadow manifolds [38]. Shadow manifolds are the result of projecting a system's signals onto coordinate axes that are time-lagged. The basis is that if a signal causes another signal, then the signal is recorded in the other signal. To prove that the control voltage (X) is actually acting on the discharge current (Y), we ran an eCCM simulation and obtained the result shown in Figure 23.

This plot shows three identifiable pairs of peaks. The orange peaks represent the discharge current causing the control voltage while the blue peaks represent the control voltage causing the discharge current. Since the first and highest orange curve peak occurs at positive time shifts, this means that discharge current strongly causes the control voltage. This makes sense as the PID algorithm directly computes setpoints based on discharge current measurements. There is a second pair of peaks after the first pair, where the order of the peaks is reversed, and where the second blue peak occurs at positive time shifts. This indicates that there is a weaker causality between control voltage and discharge current. Similarly, the third pair of peaks shows that discharge current causes control voltage. In summary, this scenario occurs where oscillation reduction is smaller.

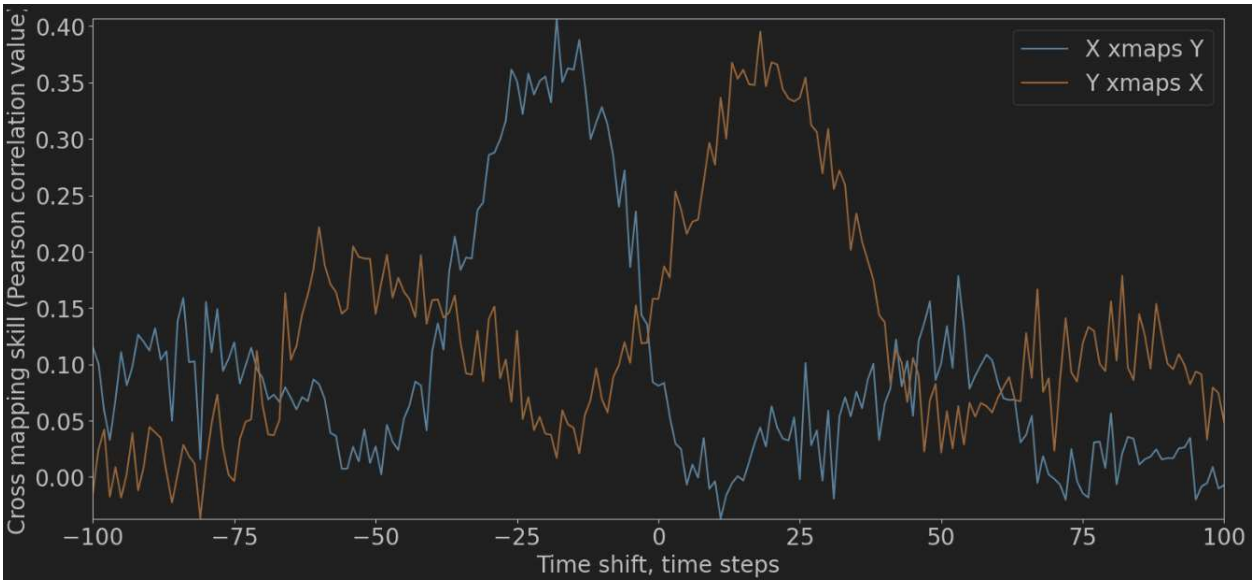
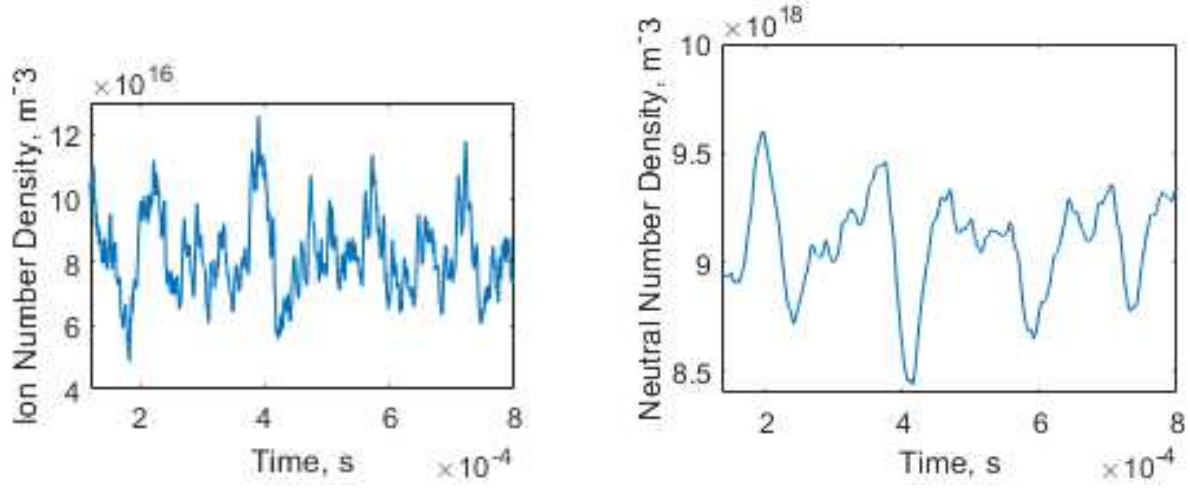


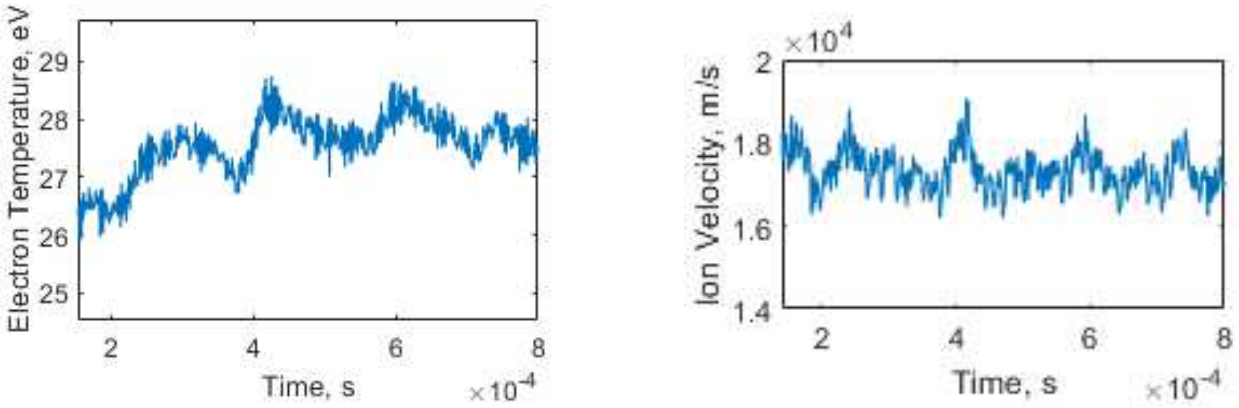
Fig. 23 eCCM Correlation and Causation of Control Voltage (X) and Discharge Current (Y) at 400 V Operating Point

EKF state estimation was used to estimate the plasma states using the discharge current measurements obtained from the PID experiment. EKF uses the 0D model presented in the Appendix and linearizes the nonlinear dynamics while predicting and correcting using the measurements [11]. In Figures 24-26, we can see that the ion number density, ion velocity, and electron temperature have major peaks and troughs that are in phase, whereas the neutral number density and electron velocity have peaks and troughs that are 180 degrees out of phase. This finding makes sense physically, as ion and neutral populations reach their maximum and minimum at the same time by way of the predator-prey interaction. The electron temperature also contributes to ionization; thus, it makes sense that the ion number density reaches a maximum when the electron temperature reaches a maximum. The ion velocity is also consistent with the ion number density, with a small ion transit time. The electron velocity and neutral number density are in phase since, as the neutral number density increases, more electrons are needed for ionization. The FFTs of the plasma parameters all exhibit the characteristic BM peak.



(a) Ion Number Density Estimates **(b) Neutral Number Density Estimates**

Fig. 24 PID EKF Plasma State Estimates (Ion and Neutral Number Densities) at 400 V Operating Point



(a) Electron Temperature Estimates **(b) Ion Velocity Estimates**

Fig. 25 PID EKF Plasma State Estimates (Electron Temperature and Ion Velocity) at 400 V Operating Point

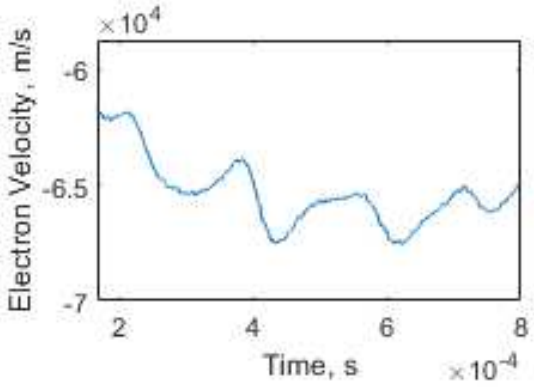


Fig. 26 PID EKF Plasma State Estimate for Electron Velocity at 400 V Operating Point

Table 9 shows the RMS of the plasma parameters with and without PID control. The reduced discharge current oscillations from using PID result in lower plasma parameter oscillations. The ion number density oscillations decrease by 7.48%, however, the electron velocity oscillations increase by 4.44%. It is possible PID excited higher frequency electron dynamics based on the fact that electrons respond almost immediately to perturbations.

Table 9 PID Plasma State Estimates at 400 V Operating Point

Plasma Parameter	Oscillation RMS (No Control)	Oscillation RMS (PID)	Oscillation RMS (% Change)
N_i	$1.176 \times 10^{16} \pm 8.124 \times 10^{12}$ m^{-3}	$1.088 \times 10^{16} \pm 7.438 \times 10^{12}$ m^{-3}	-7.48 ± 0.07
N_n	$2.185 \times 10^{17} \pm 2.427 \times 10^{13}$ m^{-3}	$2.038 \times 10^{17} \pm 2.194 \times 10^{13}$ m^{-3}	-6.74 ± 0.12
T_e	$0.5473 \pm 9.599 \times 10^{-5}$ eV	$0.508 \pm 8.692 \times 10^{-5}$ eV	-7.17 ± 1.95
U_i	429.206 ± 0.514 m/s	413.5 ± 0.484 m/s	-3.66 ± 0.16
U_e	$2.074 \times 10^3 \pm 0.145$ m/s	$2.166 \times 10^3 \pm 0.136$ m/s	4.44 ± 0.41

C. ESN Training and Testing Results at 400 V Operating Point

Using a dataset of 180 perturbation conditions consisting of control voltage and discharge current time series, we trained and tested an ESN using a single-step approach and k-fold validation with a Huber loss function. Single-step uses ground-truth data from measurements for predicting the next step, known as teacher-forced training, whereas multi-step predicts over several steps using measurements from one step. We choose single-step over multi-step for simplicity with shorter horizons in NMPC to reduce compounding error. We allow the ESN to washout before training where we eliminate the influence of the ESN's initial random state, allowing it to warm up and settle. Optuna, a Bayesian optimization framework, was used to tune the hyperparameters of the ESN within a search range, shown in Table 10, which includes the spectral radius, leak rate, sparsity, ridge alpha, reservoir size, Eta0, input scale, and feedback scale [39]. The training, testing, and tuning process took approximately one hour to complete.

Table 10 ESN Hyperparameters at 400 V Operating Point

Hyperparameter	Description	Search Range	Optimal Value
Reservoir Size	Number of neurons in the reservoir. Larger reservoirs capture temporal dependencies and nonlinearity better but if too large, then increase in computation time and risk of overfitting	[512, 1024]	960 (High)
Spectral Radius	Largest absolute eigenvalue of internal reservoir weight matrix. Controls temporal dependencies, memory strength, and stability	[0.35, 0.95]	0.877 (High)
Leak Rate	Update speed of reservoir states. Governs how fast reservoir neurons update relative to input signal	[0.1, 0.8]	0.724 (High)
Sparsity	Reservoir connectivity. Fraction of zero connections in the reservoir weight matrix	[0.988, 0.997]	0.991 (High)
Ridge Alpha	Output regularization. For ridge regression when training output weight matrix, smaller value means better fitting of data	$[1 \times 10^{-4}, 5 \times 10^{-2}]$	3.64×10^{-4} (Low)
Eta0	Initial learning rate. The step size for iterative solvers when updating weights. Larger value means faster convergence	$[8 \times 10^{-4}, 4 \times 10^{-3}]$	2.77×10^{-3} (High)
Input Scale	Input weight scaling. Scaling factor for weights from inputs to reservoir neurons. Lower value prevents input saturation	[0.3, 5.0]	1.358 (Low)
Feedback Scale	Feedback and memory strength. Scaling factor for feedback weights connecting outputs back to reservoir	[0.3, 3.0]	2.5 (High)

The R^2 value was 0.902, suggesting a strong fit by variance to the dataset, while the RMSE was 0.1917, also suggesting lower error. However, the NRMSE was 0.3122, an acceptable result but suggests that the ESN predicts the dynamics well overall, but doesn't fully capture the amplitude of the time series. This may indicate a possible limitation

in modeling HET time series data due to its NLTV and hysteretic nature. In Figure 27, we observe excellent agreement in predicting the discharge current when no control is applied for the 8-ms duration signal, with slight under prediction on some peaks.

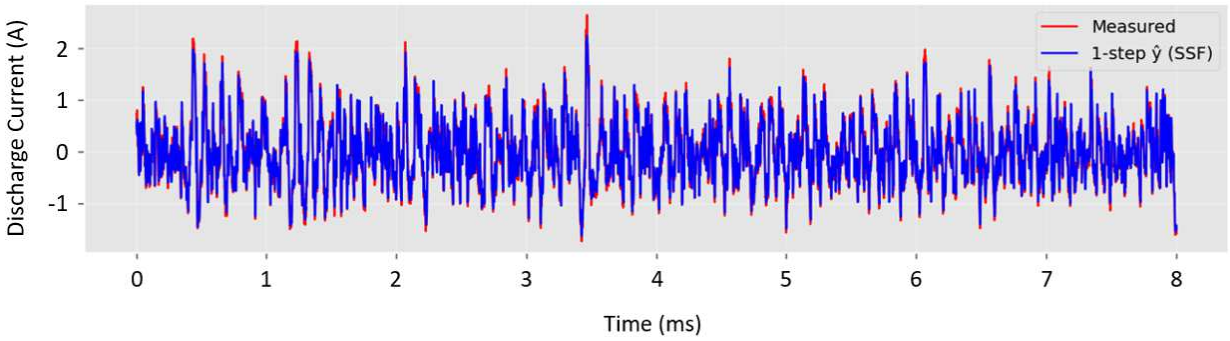
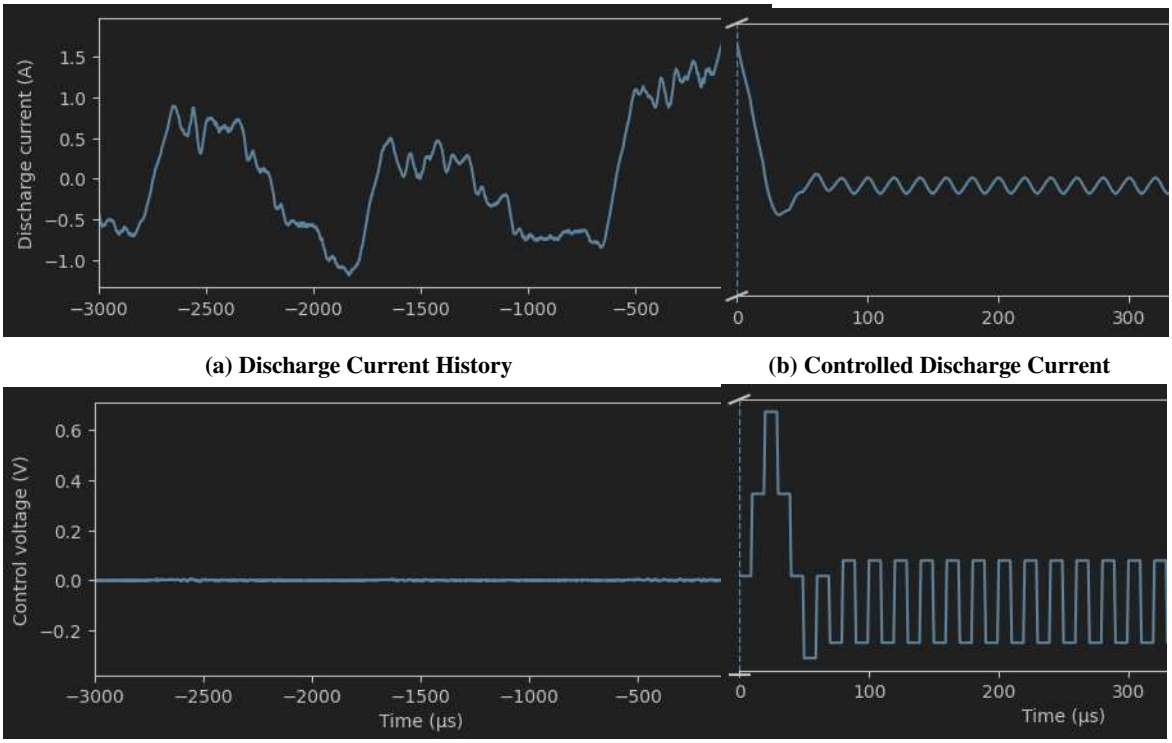


Fig. 27 Time Series Validation at 400 V Operating Point

D. ESN Model PID Simulation Results at 400 V Operating Point

In Figure 28, we see how PID affects the discharge current in an ESN simulation at a K_p value of 5 and a K_d value of 0.01, taking a few seconds to complete. We use the history of discharge current at no control to initialize the ESN, then apply PID on the ESN to minimize discharge current oscillations. We can see that the oscillations are reduced within 60 microseconds of applying controls. Figure 28d shows the chosen control values at each 10 microsecond step. The PP of the discharge current oscillations reduces from 3 A_{PP} to 0.2 A_{PP}, which is a 93.333% reduction. The control voltage ranges from -0.3 V to 0.6 V, with a close to zero mean since the injection path only takes AC signals. Accounting for gains from the amplifiers between the DAC output to the secondaries of the transformers, the perturbation voltage that would arrive at the HET ranges between -5.625 V and 11.25 V, which is 16.875 V_{PP}. Since the controlled oscillations do not match between PID experiment and PID simulation of an ESN, it is likely that the ESN has not yet fully learned the HET dynamics and that more learning iterations are needed.



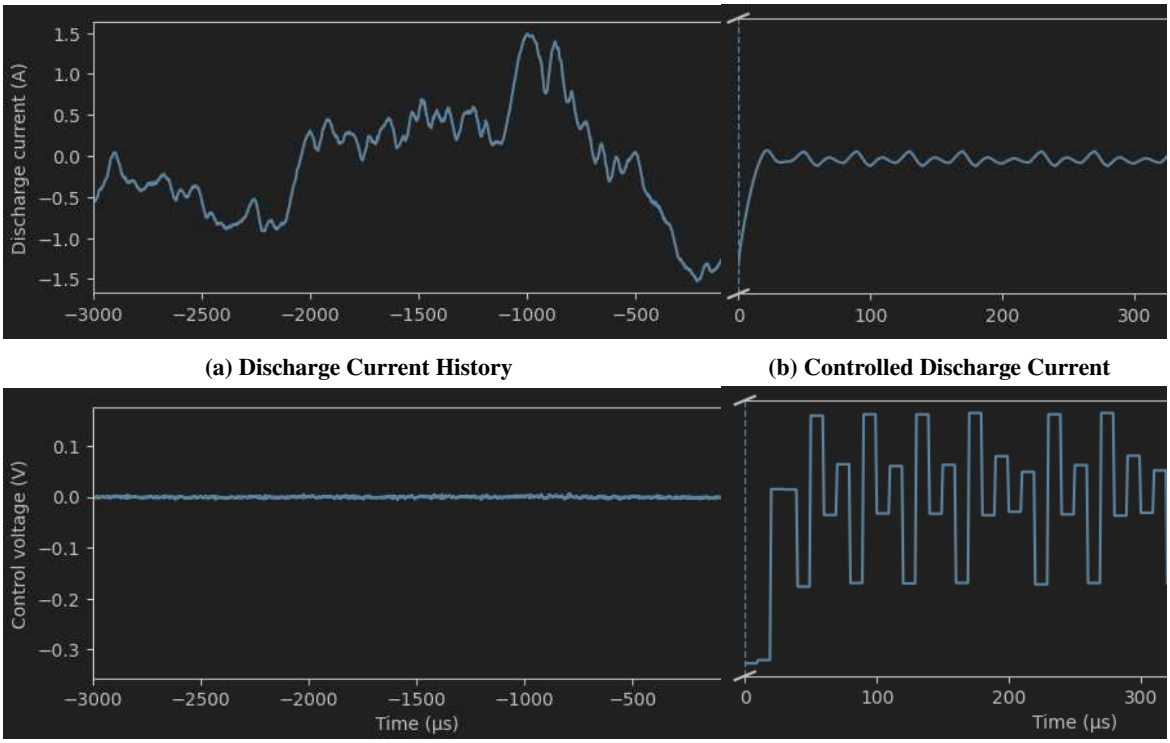
(a) Discharge Current History **(b) Controlled Discharge Current**

(c) Control History **(d) Controls**

Fig. 28 PID Control of Discharge Plasma ESN at 400 V Operating Point

E. ESN Model ML Control Simulation Results

In Figure 29, we see how NMPC affects the discharge current in an ESN simulation, which takes under a minute to complete. We use history of the discharge current at no control to initialize the ESN, and apply NMPC on the ESN to minimize discharge current oscillations. We can see that the oscillations are reduced within 30 μs of applying controls. Figure 29d shows the chosen control values at each 10 μs step. The PP of the discharge current oscillations reduces from 3 A_{PP} to 0.2 A_{PP}, which is a 93.333% reduction. The control voltage ranges from -0.3 V to 0.15 V, with zero mean, since the injection path only takes AC signals. Accounting for gains from the amplifiers between the DAC output to the secondaries of the transformers, the perturbation voltage that would arrive at the HET ranges between -5.625 V and 2.813 V, which is 8.438 V_{PP}. These results show that NMPC on an ESN can be done in simulation and could potentially be used in an experiment. The performance of PID and NMPC in terms of oscillation reduction is comparable, but NMPC requires a lower PP voltage to achieve the same result. NMPC control also acts faster than PID. The results likely will differ in an experiment as simulations do not take into account hardware delays, while repeated updating of the ESN model may reveal superior performance with ML control compared to PID control.



(a) Discharge Current History **(b) Controlled Discharge Current**

(c) Control History **(d) Controls**

Fig. 29 NMPC Control of Discharge Plasma ESN at 400 V Operating Point

VI. ML Control Experiment at 250 V Operating Point

A. Methods

Figure 30 shows the ML control experiment setup, where an FPGA is now used rather than a microcontroller for active control. We plan to use the ZCU104 FPGA, the THS1206A ADC, and the DC2459A DAC. These devices are sufficient to enable real-time control using neural networks [40]. The hardware additions could be placed within the PPU itself, but one question to be answered is whether or not the benefits of ML control outweigh the design changes. It is also possible to investigate how ML control affects the plasma using the 0D model along with state estimation using an EKF. Thrust measurements can also be done when phase synchronizing voltage and current oscillations using the ML controller.

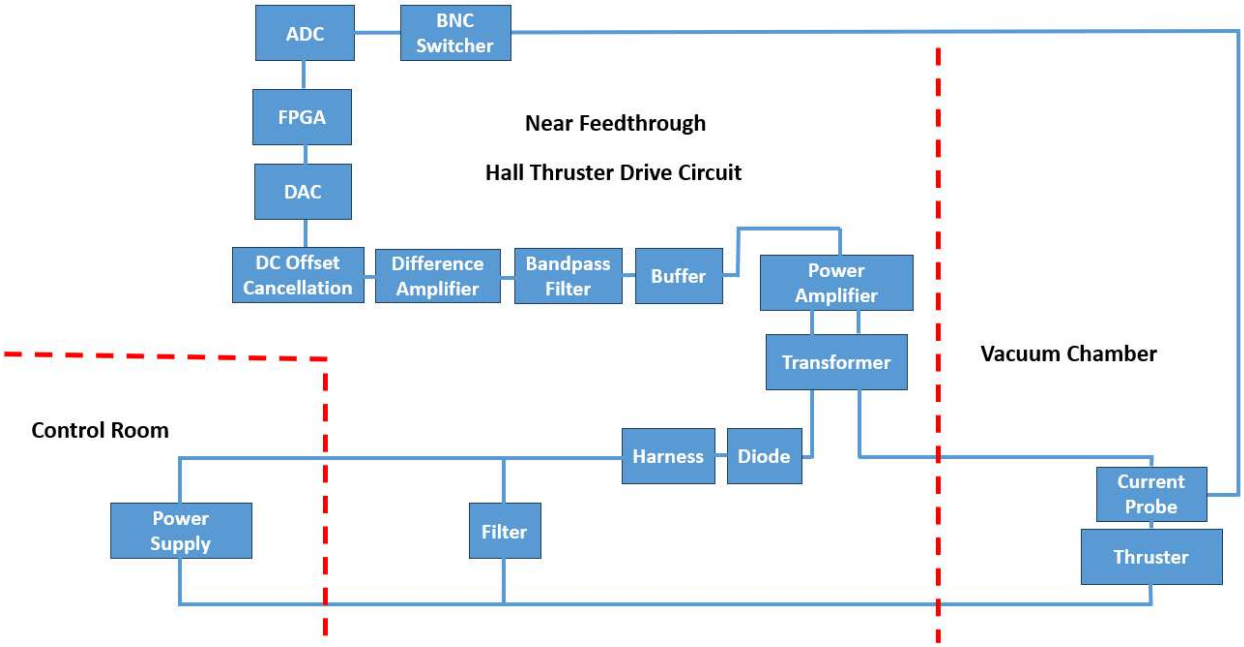


Fig. 30 Experimental setup for ML control of the HET load

The experimental flow, outlined in Figure 31, is to take the NMPC simulation results and train, test, and tune an MLP control policy using PyTorch and an NVIDIA GeForce RTX 3050 GPU on various trajectories produced by the NMPC simulation that generalize control decisions by initializing the discharge current with various time histories [34, 41]. This process is a form of imitation learning [42]. We use the MLP as it runs faster on an FPGA unlike RNNs, containing only an input layer, hidden layers, and an output layer [43]. We then program the MLP onto the FPGA using hls4ml, a HLS library for converting Python code into FPGA code, and Vitis and Vivado for microsecond inferencing [40]. The final step is to test the FPGA controller with SPSA tuning of the weights of the outer layer of the MLP to account for variations at the millisecond time scale, along with appropriate safety limits [44, 45]. Safety limits include understanding the resulting additional discharge current from applied control voltages and adding these contributions at every perturbation frequency of the total control voltage as an upper bound to define a constraint in the NMPC controller. The safety limits are tied to the impedance of the plasma in the small-signal case, as the control voltage is related to the discharge voltage, and is a transfer function that has control implications, where the plasma may be treated as LTV with small perturbations. ML control is still advantageous since some operating points may require nonlinear control to be applied with larger perturbations to better minimize oscillations. The time-varying nature of the load and the control flexibility that ML control brings are also important reasons for its use. Upon collecting time series data with the controller on, the process flow can be repeated by updating the ESN, potentially online, in under a minute to then update the MLP control policy, resulting in iterative improvement of the controller with subsequent testing.

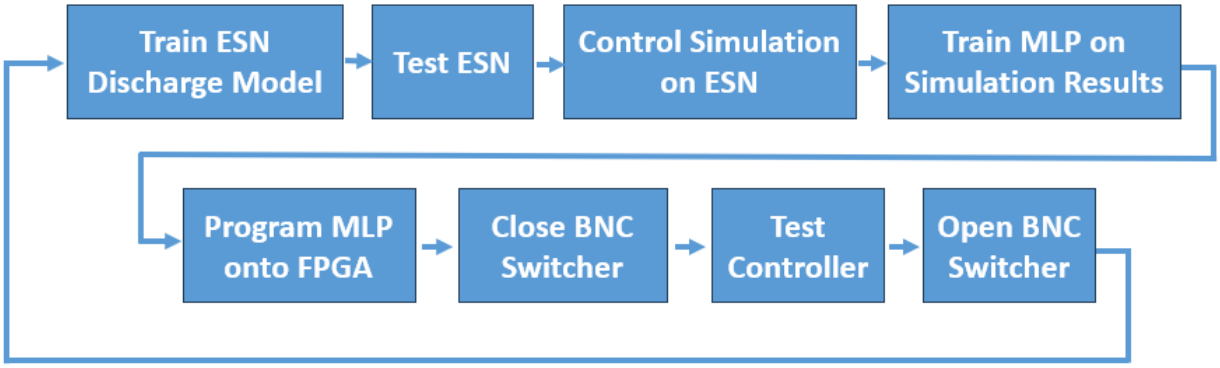


Fig. 31 ML Control Testing Process

Figure 37 shows the experimental setup used to verify that a neural network is properly implemented on an FPGA, by comparing the ground truth NN output from various excitations in simulation to those obtained from experiment. This setup represents a hardware-in-the-loop strategy for validating NNs.

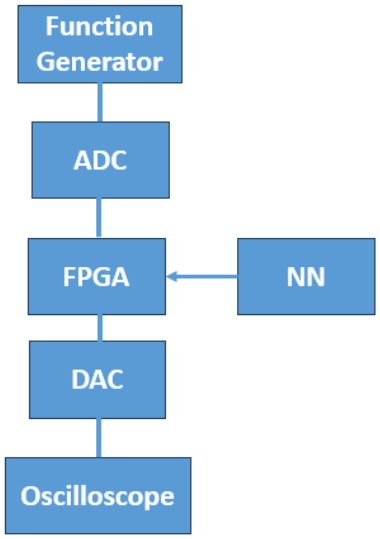


Fig. 32 FPGA NN Open Loop Testing

B. Results

Approximately 20 minutes of simulation time was taken to generate NMPC control trajectories from various discharge current waveform histories, creating a control dataset of 266 trajectories, with a total of 18620 training samples. The MLP was trained using gradient descent on this dataset in under one minute of runtime. The MLP consists of an input layer, two hidden layers, and an output layer, represented by the architecture 40 -> 128 -> 64 -> 1, denoting 40 neurons in the input layer, 128 and 64 neurons in the first and second hidden layers, and lastly one neuron in the output layer.

Figure 33 shows the result of testing the MLP controller against actual and previously unseen control trajectories from the dataset. An R2 value of 0.8577, an RMSE of 0.243, and an NRMSE of 0.3772 were achieved, demonstrating sufficient prediction accuracy for use in a HET experiment. It can be observed that the predicted controls roughly follow the unseen control trajectory from the dataset.

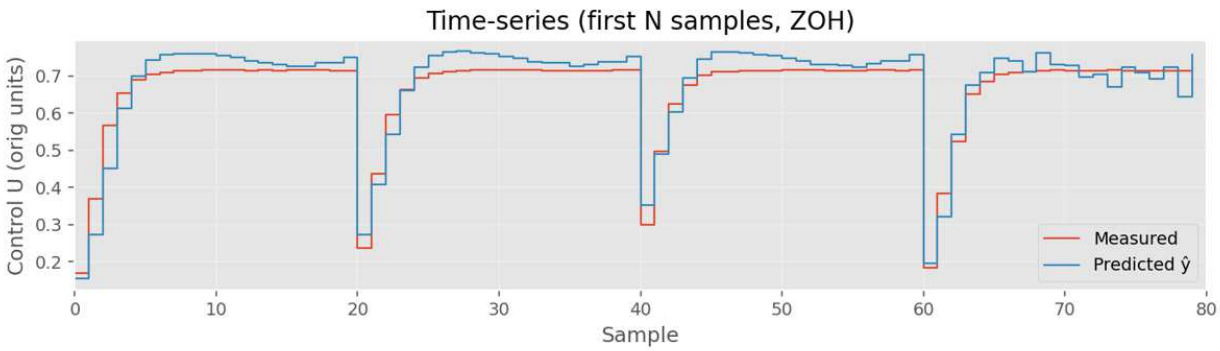


Fig. 33 MLP Controller Testing

Figure 34 shows the latency of the MLP implemented on the FPGA, along with layer precisions tuned using the C simulator within Vitis, layer activation functions, the type of pipelining, BRAM usage, DSP usage, the number of flip-flops, and lastly the number of look-up tables. A total latency of 785 ns exists for neural network inference to determine control outputs from incoming discharge current measurements, which meets the control loop budget, maximizing the amount of time a control voltage setpoint can be held within each 10 microsecond period.

MODULES & LOOPS	LATENCY(NS)	ITERATION LATENCY	INTERVAL	TRIP COUNT	PIPELINED	STRUCTURE	BRAM	DSP	FF	LUT
> myproject (6)	785.000		64		dataflow	function	184	258	59,363	63,284
> dense_resource_rf_leq_nin_ap_fixed_ap_fixed_12_4_5_3_0_config2_s (1)	210.000		40		loop rewind st	function	43	128	7,850	7,723
tanh_ap_fixed_12_4_5_3_0_ap_fixed_12_4_5_3_0_tanh_config3_s	10.000		1		yes	function	64	0	2,820	12,943
> dense_resource_rf_leq_nin_ap_fixed_ap_fixed_12_4_5_3_0_config4_s (1)	340.000		64		loop rewind st	function	43	128	7,428	7,987
tanh_ap_fixed_12_4_5_3_0_ap_fixed_12_4_5_3_0_tanh_config5_s	10.000		1		yes	function	32	0	1,412	6,479
> dense_resource_rf_leq_nin_ap_fixed_ap_fixed_16_1_5_3_0_config6_s (1)	180.000		32		loop rewind st	function	1	2	1,711	2,226
tanh_ap_fixed_16_1_5_3_0_ap_fixed_16_1_5_3_0_tanh_config7_s	10.000		1		yes	function	1	0	23	77

Fig. 34 Neural Network Latency on FPGA

Figure 35 shows the FPGA usage with implemented areas in light blue representing DSP and dark blue representing BRAM. There is still additonial space for more complex designs seeing that the latency is only 785 ns. More

representative neural networks could be used for greater performance.

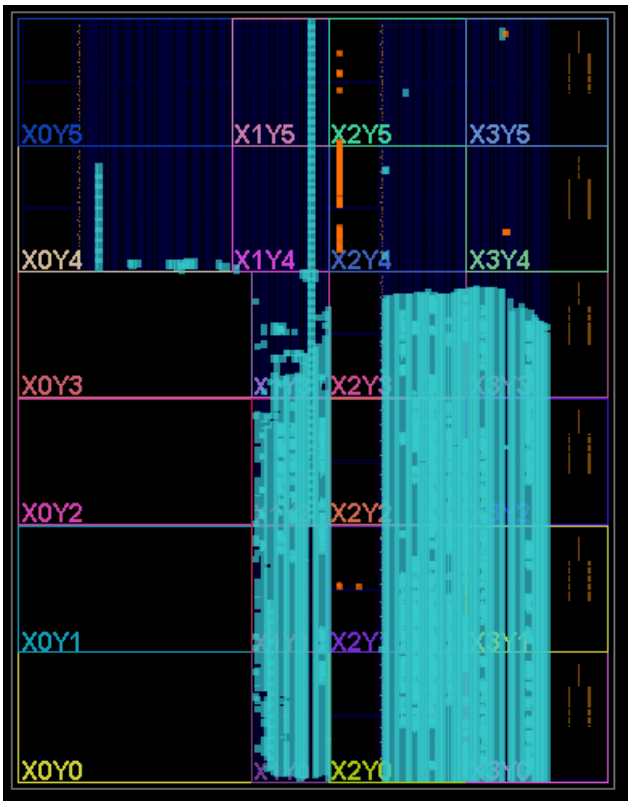


Fig. 35 FPGA Usage

Figures 36 and 37 show simulation results and experimental results of a MLP perturbed by a 1 A amplitude and 1 kHz sinusoid. It can be observed from the two figures that there is agreement between the experimental result and the simulation result. This shows that the MLP was successfully implemented on the FPGA, making it ready for use in a HET experiment.

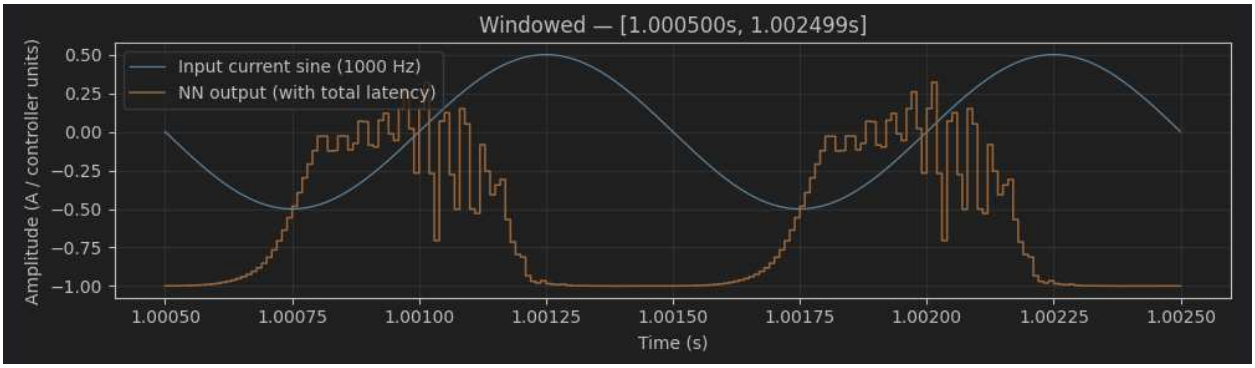


Fig. 36 FPGA NN Simulation Results

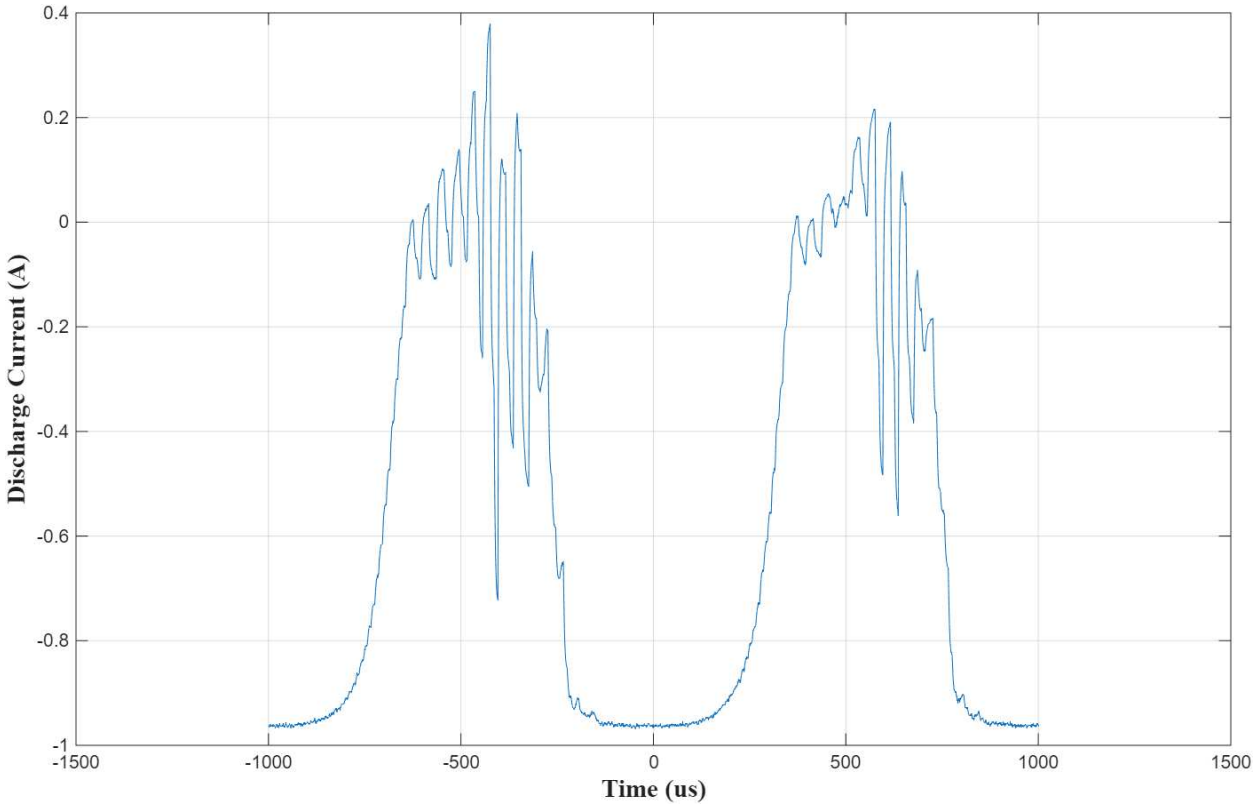


Fig. 37 FPGA NN Open Loop Experiment Results

Figure 38 shows how discharge current changes with each iteration of the SPSA algorithm. Upon each iteration, the weights are adjusted in an attempt to reduce the RMS of the discharge current. We can see that the algorithm is not completely stable and not fully optimized, but at iteration 48, the discharge current reaches a minimum of 2.7195 A, which is a 1.33% reduction from the unperturbed mean discharge current of 2.7563 A. We determined that this percentage reduction is not due primarily due to noise, as the variability in the RMS from one block of 1 million samples to another leads to variation less than 10 mA, whereas the reduction is by 36.8 mA. This reduction outperforms PID at the same operating point, which increased the oscillations by 4.17%.

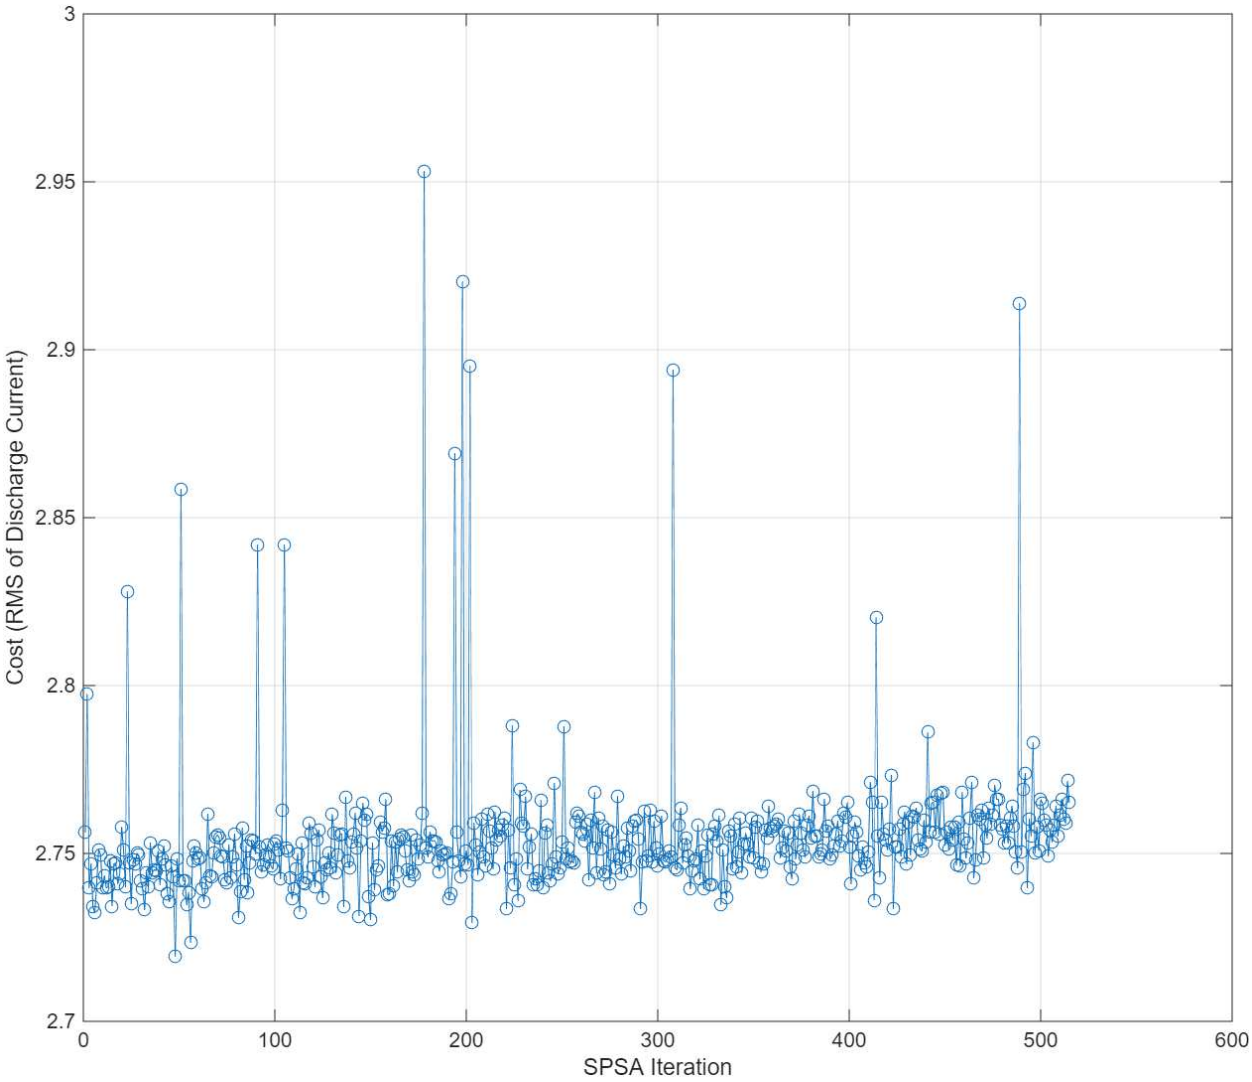


Fig. 38 SPSA RMS vs Iteration

Figure 39 shows how the 64 weights changed after each SPSA iteration. One detail that can be observed is that many of the weights start farther apart, but as the iteration count grows, the weights become closer in value and start to stabilize. During these iterations, it is possible to revert to the best combination of weights determined, to operate with the best oscillation reduction observed.

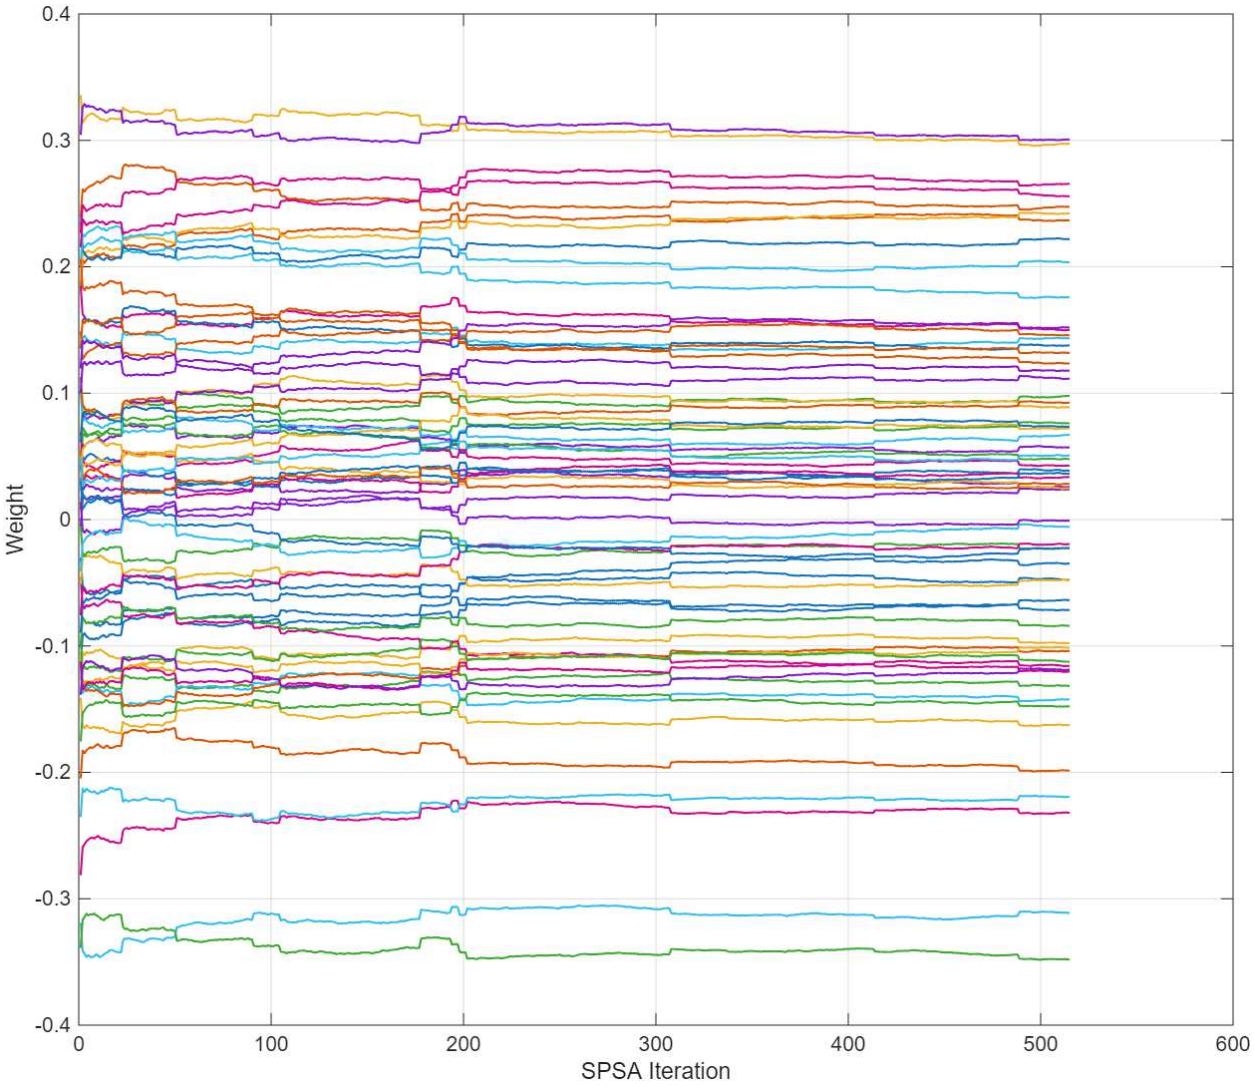


Fig. 39 SPSA Weights vs Time

The results of the ML control experiment show that it is indeed possible to use this approach to control oscillations better than PID, which increased rather than decreased oscillations by 4.17 % at the 250 V operating point. Some reasons for why the control extent was limited is driven by the operating point, but to a much larger extent that the perturbation data set was limited, and training was not performed online to account for changes in the plasma. An online adaptive approach would allow for a better discharge plasma model that could lead to better control trajectories and reduced oscillations.

VII. Conclusion

The research described here attempts to create an intelligent power supply that can control HETs in an optimal manner. We show that perturbations, in the form of PID control, reduce oscillations and impact the discharge current using the eCCM algorithm, revealing reciprocal causality. We use the EKF to determine how plasma parameters change with and without PID control. We demonstrate the effectiveness of ESNs to predict time-resolved discharge current data from a HET by using a diverse dataset of perturbation response data, yielding 90% prediction accuracy by variance. We show that a real-time ML control approach can potentially minimize oscillations through simulation, reducing oscillations by more than 90%. We then use the outputs of the ML control simulations to train a MLP that can make control predictions, which is then deployed to an FPGA using hls4ml, Vitis, and Vivado. We demonstrate the controller through an experiment that successfully shows reduced discharge current oscillations in comparison to PID control.

The ML controller has the potential to work at other operating points, as the same procedure of learning dynamics and computing controls does not change. With advances in rad-hardened GPU and FPGA technology, it has become increasingly possible to employ ML in space missions.

Declarations

A. Funding

This work was sponsored by the NASA Joint Advanced Propulsion Institute under the grant number 80NSSC21K1118.

B. Conflicts of Interest/Competing Interests

The authors have no competing interests to declare that are relevant to the content of this article.

C. Data Availability

N/A

D. Code Availability

N/A

E. Author's Contributions

A. Krishnan developed the concept, executed the study, and wrote the paper. B. Vlahov, J. Gibson, and E. Theodorou helped develop the machine learning control framework used in the paper. D. Lev and M. Walker provided guidance on the plasma physics parts of the study. M. Saeedifard provided guidance on the electrial engineering parts of the study.

Acknowledgments

The authors would like to thank all of the undergraduate and graduate student researchers at HPEPL that contributed towards experimental setup and vacuum facility operation, particularly Julian Lopez-Uricoechea, Will Brabston, Nafez Ahmmed, Frank George, Hector Sanchez, Irene Moxley, and Daniel Wochnick. The authors would also like to thank the Georgia Tech Research Corporation for assisting in filing a patent application, US 63/765,398, based on the work presented in this conference paper.

References

[1] Brown, D. L., "Investigation of low discharge voltage Hall thruster operating modes and ionization processes," *31st International Electric Propulsion Conference*, Ann Arbor, MI, USA, 2009. IEPC-2009-074.

[2] Goebel, D., and Katz, I., *Fundamentals of Electric Propulsion: Ion and Hall Thrusters*, Wiley, Hoboken, NJ, 2008.

[3] Byers, D., "A review of facility effects on Hall effect thrusters," *31st International Electric Propulsion Conference*, Ann Arbor, MI, USA, 2009. IEPC-2009-076.

[4] Frieman, J. D., "Role of a conducting vacuum chamber in the Hall Effect Thruster Electrical Circuit," *Journal of Propulsion and Power*, 2014.

[5] Byrne, M. P., "Coupling of electrical and pressure facility effects in Hall effect thruster testing," *37th International Electric Propulsion Conference*, Boston, MA, USA, 2022. IEPC-2022-377.

[6] Walker, J. A., "Electrical facility effects on hall current thrusters: Electron termination pathway manipulation," *Journal of Propulsion and Power*, 2016.

[7] Dankanich, J. W., "Recommended practice for pressure measurement and calculation of effective pumping speed in electric propulsion testing," *Journal of Propulsion and Power*, 2017.

[8] Choueiri, E. Y., "Plasma oscillations in hall thrusters," *Physics of Plasmas*, 2001.

[9] Fife, J., Martinez-Sanchez, M., Szabo, J., Fife, J., Martinez-Sanchez, M., and Szabo, J., "A numerical study of low-frequency discharge oscillations in Hall thrusters," *33rd Joint Propulsion Conference and Exhibit*, 1997.

[10] Barral, S., and Ahedo, E., "Theoretical study of the breathing mode in Hall thrusters," *42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, 2006.

[11] Troyetsky, D. E., Greve, C. M., Tsikata, S., and Hara, K., "State estimation for real-time analysis of dynamic plasma properties and electrical circuit effects in Hall effect thrusters," *37th International Electric Propulsion Conference*, Boston, MA, USA, 2022. IEPC-2022-323.

[12] Barral, S., "Closed-loop control of ionization oscillations in Hall accelerators," *Physics of Plasmas*, 2011.

[13] Tamida, T., "Performance Improvement of Hall Thrusters Using a Pulse-Synchronous Driver System," *Journal of Propulsion and Power*, 2015.

[14] Li, W., "Effect of Low-frequency Oscillation on Performance of Hall Thrusters," *35th International Electric Propulsion Conference*, Atlanta, GA, USA, 2017. IEPC-2017-134.

[15] Petronio, F., "Study of the breathing mode development in Hall thrusters using hybrid simulations," *Journal of Applied Physics*, 2024.

[16] Krishnan, A., "Impedance Analysis of the Hall Thruster Discharge Circuit and Plasma Load to Address Harness Facility Effects," *38th International Electric Propulsion Conference*, Toulouse, France, 2024. IEPC-2024-619.

[17] Ito, K., "Sequential prediction of Hall thruster performance using echo state network models," *Transactions of the Japan Society for Aeronautical and Space Sciences*, 2024.

[18] Jovel, D. R., "Hall effect thruster impedance characterization in ground-based vacuum test facilities," *Journal of Electric Propulsion*, 2024.

[19] Thoreau, P., "Rapid thruster-in-the-loop optimization for Hall thrusters," *Journal of Electric Propulsion*, 2025.

[20] Barral, S., "Numerical investigation of closed-loop control for Hall accelerators," *Journal of Applied Physics*, 2011.

[21] Wei, L., "Stabilizing Low Frequency Oscillation with Two Stages Filter in Hall Thrusters," *34th International Electric Propulsion Conference*, Shanghai, China, 2015. IEPC-2015-284.

[22] Brooks, J., "Iterative learning control of the hollow cathode plume mode," *38th International Electric Propulsion Conference*, Toulouse, France, 2024. IEPC-2024-219.

[23] Simmonds, J., "Studies of a modulated hall thruster," *Plasma Sources Science and Technology*, 2021.

[24] Slimane, T. B., "Analysis and control of hall effect thruster using optical emission spectroscopy and artificial neural network," *Journal of Applied Physics*, 2024.

[25] Shaoning, W., "Aerospace Digital and Intelligent Power Supply Technology Research," *FAFEE*, 2024.

[26] Hopfield, J. J., "Neural networks and physical systems with emergent collective computational abilities," *PNAS*, 1982.

[27] Sekarak, M. J., "Mode Transitions in Hall Effect Thrusters," *49th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, Huntsville, AL, USA, 2013. AIAA-2013-4116.

[28] Piñero, L. R., "The Impact of Harness Impedance on Hall Thruster Discharge Oscillations," *NASA/TM-2017-219722*, 2017.

[29] Priese, S., "Adaptive feedforward control for active noise cancellation in-ear headphones," *Proceedings of Meetings on Acoustics*, 2012.

[30] Langendorf, S., "Effects of wall electrodes on Hall effect thruster plasma," *Physics of Plasmas*, 2015.

[31] Jorns, B. A., "A Predictive Hall Thruster Model Enabled by Data-Driven Closure," *AIAA Propulsion and Energy Forum*, 2020.

[32] Armenio, L. B., "Echo State Networks: analysis, training and predictive control," *18th European Control Conference*, Napoli, Italy, 2019.

[33] Jordanou, J. P., "Echo state networks for practical nonlinear model predictive control of unknown dynamic systems," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.

[34] Popescu, M.-C., "Multilayer perceptron and neural networks," *WSEAS Transactions on Circuits and Systems*, 2009.

[35] Kieckhafer, A. W., "Recirculating liquid nitrogen system for operation of cryogenic pumps," *32nd International Electric Propulsion Conference*, Kurhaus, Wiesbaden, Germany, 2011. IEPC-2011-217.

[36] Wachter, A., "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, 2005.

[37] Andersson, J. A. E., "CasADi: a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, 2018.

[38] Huerta, C. E., "Determining causality in hall effect thrusters using extended convergent cross mapping, part I," *Plasma Sources Science and Technology*, 2021.

[39] Akiba, T., "Optuna: A Next-generation hyperparameter optimization framework," *KDD*, Anchorage, AK, USA, 2019.

[40] Shi, R., "ML-based Real-Time Control at the Edge: An Approach Using hls4ml," *IEEE International Parallel and Distributed Processing Symposium Workshop*, San Francisco, CA, USA, 2024.

[41] Paszke, A., "PyTorch: an imperative style, high-performance deep learning library," *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Vancouver, Canada, 2019.

[42] Zare, M., "A Survey of Imitation Learning: Algorithms, Recent Developments, and Challenges," *IEEE Transactions on Cybernetics*, 2024.

[43] Vaithianathan, M., "Comparative Study of FPGA and GPU for High-Performance Computing and AI," *ESP International Journal of Advancements in Computational Technology*, 2023.

[44] Spall, J. C., "Implementation of the Simultaneous Perturbation Algorithm for Stochastic Optimization," *IEEE Transactions on Aerospace and Electronic Systems*, 1998.

[45] Hewing, L., "Learning-based model predictive control: toward safe learning in control," *Annual Review of Control, Robotics, and Autonomous Systems*, 2020.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- MachineLearningControlJEP2026Appendix.pdf